Efficient and Flexible Crowdsourcing of Specialized Tasks with Precedence Constraints

Avhishek Chatterjee, Michael Borokhovich, Lav R. Varshney, and Sriram Vishwanath

arXiv:1601.04094v1 [cs.MA] 15 Jan 2016

Abstract

Many companies now use crowdsourcing to leverage external (as well as internal) crowds to perform specialized work, and so methods of improving efficiency are critical. Tasks in crowdsourcing systems with specialized work have multiple steps and each step requires multiple skills. Steps may have different flexibilities in terms of obtaining service from one or multiple agents, due to varying levels of dependency among parts of steps. Steps of a task may have precedence constraints among them. Moreover, there are variations in loads of different types of tasks requiring different skill-sets and availabilities of different types of agents with different skill-sets. Considering these constraints together necessitates the design of novel schemes to allocate steps to agents. In addition, large crowdsourcing systems require allocation schemes that are simple, fast, decentralized and offer customers (task requesters) the freedom to choose agents. In this work we study the performance limits of such crowdsourcing systems and propose efficient allocation schemes that provably meet the performance limits under these additional requirements. We demonstrate our algorithms on data from a crowdsourcing platform run by a non-profit company and show significant improvements over current practice.

I. Introduction

The nature of knowledge work has changed to the point nearly all large companies use *crowdsourcing* approaches, at least to some extent [2]. The idea is to draw on the cognitive energy of people, either within a company or outside of it [3]. A particularly notable example is the non-profit *impact sourcing* service provider, Samasource, which relies on a marginalized population of workers to execute work, operating under the notion *give work, not aid* [4], [5].

There are multifarious crowdsourcing structures [6], [7] that each require different strategies for matching work to agents [8]. Contest-based platforms such as TopCoder and InnoCentive put out open calls for participation, and best submissions win prizes [9]. Microtask platforms such as Amazon Mechanical Turk allocate simple tasks on a first-come-first-serve basis to any available crowd agent. When considering platforms with skilled crowds and

A. Chatterjee and L. R. Varshney are with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, Illinois, USA. (email: {avhishek,varshney}@illinois.edu).

M. Borokhovich is with the AT&T Labs, New Jersey, USA. (email: michaelbor@utexas.edu).

S. Vishwanath is with the Wireless Networking and Communication Group, The University of Texas at Austin, Austin, Texas, USA. (email: sriram@austin.utexas.edu).

Part of the material in this paper will be presented at IEEE INFOCOM 2016, San Francisco, USA [1].

specialized work, such as oDesk (now upWork) [7], IBM's Application Assembly Optimization platform [10], and to a certain extent Samasource's SamaHub platform [4], efficient allocation algorithms are needed.

In these skill-based crowdsourcing platforms, the specialized tasks have multiple steps, each requiring one or more skills. For example, software development tasks may first be planned (architecture), then developed (programming), and finally tested (testing and quality assurance), perhaps with several iterations. Even in skilled microtasking platforms like SamaHub, most jobs have more than one step. Task steps often have *precedence constraints* between them, implying that a particular step of a task can only be performed after another set of steps has been completed.

To serve a step requiring multiple skills, we need either a single agent that has all of the skills or a group of agents that collectively do so. Whether multiple agents can be pooled to serve a step or not depends on the *flexibility* of the step: if there are strong interdependencies between different parts of a step, the step may require a single agent. Notions of flexibility and precedence constraints are central to this paper.

Allocating tasks to servers is a central problem in computer science [11], communication networks [12], and operations research [13]. The skill-based crowdsourcing setting, however, poses new challenges for task allocation in terms of vector-valued service requirements, random and time-varying resource (agents) availability, large system size, a need for simple decentralized schemes requiring minimal actions from the platform provider, and the freedom of customers (task requesters) to choose agents without compromising system performance. Some of these issues have been addressed in recent work [14], [15], but previous work does not address precedence constraints or step flexibility. The notion of flexibility in [15] is based on agent-categories and is different from here.

Task allocation with precedence constraints has been studied in theoretical computer science, as follows. Given several tasks, precedence constraints among them, and one or more machines (either same or different speed), allocate tasks to minimize the weighted sum of completion times or maximum completion time [16]. In crowdsourcing, we have a stream of tasks arriving over time and so we are interested in dynamics.

Dynamic task allocation with precedence constraints has recently been studied in [17] for Bernoulli task arrivals. This is different from crowdsourcing scenarios, and the optimal scheme is required to search over the set of possible allocations, which is not suitable for crowdsourcing systems due to their inherent high-dimensionality (many types of tasks). Additional challenges in a crowdsourcing platform are: (i) random and time-varying agent availability; (ii) vector-valued service requirements; (iii) fast computation requirements for scaling; and (iv) freedom of choice for customers.

Here we address the above issues for various flexibilities of steps and agents, to characterize limits of crowd systems and develop optimal, computationally-efficient, centralized allocation schemes. Based on insights garnered, we further present fast decentralized greedy schemes with good performance guarantees. To complement our theoretical results, we also present numerical studies on real and synthetic data, drawn from Samasource's SamaHub platform.

The remainder of the paper is organized as follows. Sec. II describes the system model for crowdsourcing platforms with different precedence and flexibility constraints. Sec. III presents a generic characterization of the system limits and a generic centralized optimal allocation scheme. Secs. V–VII address particular systems with

different flexibility constraints to yield fast decentralized schemes that meet crowdsourcing platform requirements mentioned above. Sec. VIII presents numerical studies on real and synthetic data. Detailed proofs of theoretical results are given in Appendix A.

II. System Model

There are a total of S kinds of skills available in the crowdsourcing system, numbered $[S] = \{1, 2, \ldots, S\}$. We define types of agents by skills, and denote the total number of types of agents by M. An agent of type m has skills $S_m \subset [S]$.

Tasks posted on the platform are of N types. Each type of task j has one or multiple steps associated with it, denoted K_j. A step $k \in \{1, 2, \ldots, K_j\}$ of a job type j—a (j,k)-step—needs a skill-hour service vector $r_{j,k} \in \mathbb{R}_+^S$ (non-negative orthant), i.e. $r_{j,k,s}$ hours of skill s. A part of a step of type (j,k) involving skill s is called a (j,k,s)-substep if $r_{j,k,s} > 0$, the size of this substep.

In the platform, allocations of work to available agents happen at regular time intervals, $t = 1, 2, \ldots$. Tasks that arrive after an epoch t are considered for allocation at epoch $t+1$, based on the available agents at that epoch. Tasks or parts of tasks that remain unallocated due to insufficient available skilled agents are considered again in the next epoch. We assume that for any substep (j,s), the time requirement is less than the duration between two allocation epochs.

Tasks arrive according to a stochastic process in $\mathbb{Z}_+^N$ (non-negative orthant), $\mathbf{A}(t) = (A_1(t), A_2(t), \ldots, A_N(t))$, where $A_i(t)$ is the number of tasks of type i that arrive between epochs $t-1$ and t. The stochastic process of available agents at epoch t is $\mathbf{U}(t) = (U_1(t), U_2(t), \ldots, U_M(t))$. We assume $\mathbf{A}(t)$ and $\mathbf{U}(t)$ are independent of each other and that each of the processes are i.i.d. for each t, with bounded second moments. Let $\Gamma(\cdot)$ be the distribution function of $\mathbf{U}(t)$, and let $\boldsymbol{\lambda} = \mathbf{E}[\mathbf{A}(t)]$ and $\boldsymbol{\mu} = \mathbf{E}[\mathbf{U}(t)]$ be the means of the processes.

An agent is *inflexible* if it has pre-determined how much time to spend on each of its skills. Inflexible agents bring a vector $\mathbf{h}_m = (h_{m,1}, h_{m,2}, \ldots, h_{m,S})$ where $h_{m,s} > 0$ if and only if $s \in S_m$ and an inflexible agent spends no more than $h_{m,s}$ time for skill s. Contrariwise, *flexible* agents bring a total time h_m which can be arbitrarily split across skills in S_m.

A step is *flexible* if it can be served by any collection of agents pooling their service-times. All substeps of *inflexible* steps must be allocated to one agent. At any epoch t only an integral allocation of a step is possible. Hence, in any system for a step to be allocated, all of its substeps must be allocated.

A set of flexible substeps $sst_1, sst_2, \ldots, sst_n$ of size $x_1, x_2, \ldots, x_n$ with skill s can be allocated to agents $1, 2, \ldots, m$ if the *available* skill-hours[1] for skill s of these agents, $y_1, y_2, \ldots, y_m$, satisfy the following for some $\{v_{pq} \geq 0 : 1 \leq p \leq n, 1 \leq q \leq m\}$,

$$\sum_{p=1}^{n} v_{ip} \leq y_i \ i \in [m], \sum_{q=1}^{m} v_{qj} \geq x_j \ j \in [n], \tag{1}$$

[1]Available skill-hour is determined by the availability of the agent, system state, and whether agents are flexible or inflexible.

where the $\{v_{pq}\}$ capture how agents split their available skill hours across substeps.

For inflexible steps, a set of steps $st_1, st_2, \ldots, st_n$ of size (vectors) $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ can be allocated to an agent with available skill-hours (vector) $\mathbf{y}$ if

$$\sum_{i=1}^{n} \mathbf{x}_i \leq \mathbf{y}. \tag{2}$$

There may also be precedence constraints on the order in which different steps of a task of type j can be served. For any task of type j, this constraint is given by a directed rooted tree T_j on K_j nodes where a directed edge $(k \to k'), k, k' \in [K_j]$ implies step k' of a task of type j can only be served after step k of the same task has been completed.

Scalings of several crowd system parameters are as follows. Task arrival rate $\lambda(N) = \sum_{j=1}^{N} \lambda_j$ scales faster than number of task types N, i.e. $\lim_{N\to\infty} N/\lambda(N) = 0$. Number of skills S scales slower than N, i.e. $S = o(N)$. In practice, a task requires a constant number of skills d, which implies $\Omega(S^d)$ possible types of steps. Number of skills of an agent is $d' = O(1)$ implying $M = \sum_l M^l = O(S^{d'})$, implying $M = O(N_1^{\alpha})$ and $M = \Omega(N^{\alpha_2})$ for $0 < \alpha_2 < \alpha_1$. Further, the length of tasks and availability of agents do not scale with the size of crowdsourcing systems.

Beyond these practical system scalings, we make the following mild assumptions: $\lambda_j = \omega(1)$ for all j and $\lambda^s(N)) = \sum_{j:r_{j,1,s}>0} \lambda_j(N) = \Omega\left(N^c\right) \forall s \in [S]$, for some $c > 0$. These assumptions mean the arrival rate of every type of job and the total number of jobs requiring a particular skill scale with the system. We call these scaling patterns *crowd-scaling*.

III. Notions of Optimality

To formally characterize the maximal supportable arrival rate of tasks, we introduce some more notation and invoke some well-accepted notions used in this regard.

For each $j \in [N]$, let the number of unfinished tasks in the system *just after* allocation epoch $t-1$ be $Q_j(t)$. $A_j(t)$ is the number of tasks of type j arriving between epochs $t-1$ and t. The number of tasks of type j completely allocated (all steps) at epoch t is $D_j(t)$. Thus $Q_j(t)$ evolves as:

$$Q_j(t+1) = Q_j(t) + A_j(t) - D_j(t). \tag{3}$$

Clearly $D_j(t) \leq Q_j(t) + A_j(t)$ at any epoch t, since at most $Q_j(t) + A_j(t)$ type j tasks are available. Hence $Q_j(t) \geq 0$ for all t. Note that due to additional precedence constraints, typically $D_j(t) < Q_j(t) + A_j(t)$.

Definition 1. *A scheme of allocation of tasks is called a* policy *if it allocates tasks at a time epoch t based on knowledge of statistics of the processes $\mathbf{A}$ and $\mathbf{U}$ and their realizations up to time t, but does not depend on future values.*

Definition 2. *A crowd system is* stable *under policy $\mathcal{P}$ if the process $\mathbf{Q}(t) = (Q_j(t), j \in [N])$ has a finite expectation in steady-state under that policy, i.e., $\limsup_{t\to\infty} \mathbf{E}[Q_j(t)] < \infty$, for all j for any initial condition.*

Definition 3. *An arrival rate $\boldsymbol{\lambda}$ is* stabilizable *if there exists a policy $\mathcal{P}$ under which $\mathbf{Q}(t) = (Q_j(t), j \in [N])$ is stable.*

Definition 4. *The* capacity region *of a crowd system for a given distribution Γ of the agent-availability process $\mathbf{U}(t)$ is the closure of the set $\mathcal{C}_\Gamma = \{\boldsymbol{\lambda} : \boldsymbol{\lambda}$ is stabilizable$\}$.*

We aim to propose statistics-agnostic, computationally simple and decentralized schemes that offer customers freedom of choice while stabilizing any arrival rate in the system's capacity region. Stronger than stability, often we give high probability bounds on number of unallocated tasks.

IV. Capacity and Centralized Allocation Routine

Here we present a generic characterization of the capacity region of a crowd system for all combinations of agent- and task-flexibility. We also give a generic centralized allocation routine that can be easily adapted to a particular system.

For any given set of available agents $\mathbf{u} = (u_i : 1 \le i \le M)$, define the number of different types of steps $(\{a_{j,k}\})$ that can *potentially* be allocated in a crowd system by $C(\mathbf{u}) \subset \mathbb{R}_+^{\sum_j K_j}$. When we say $\{a_{j,k}\}$ is the number of steps of different types that can potentially be allocated, we consider the following scenario that satisfies the allocation constraints in Sec. II.

- A1 An infinite number of steps of each type (j,k), $k \in [K_j]$ for a $j \in [N]$ are available for allocation, i.e., the limitation only comes from the available resource $\mathbf{u}$.
- A2 Precedence constraints among the steps are already satisfied, i.e., all corresponding (j,k)-steps of the available (j,k')-steps have already been allocated previously. This is equivalent to an absence of precedence constraints.
- A3 Integral steps must be allocated, i.e., all substeps of a step need to be allocated for allocation of the step.
- A4 To allocate $a_{j,k}$ steps of different types to a collection of R agents of type $\{m_r : r \in [R]\}$ and available hours $\{y_{m_r,s} : r \in [R]\}$ (which is a function of $\{\mathbf{h}_{m_r} : r \in [R]\}$ depending on the system), we need to satisfy either (1) or (2) depending on system type.

Let $C^{\text{cvx}}(\mathbf{u})$ be the convex hull of the set $C(\mathbf{u})$, and define another set $\mathbf{C} \subset \mathbb{R}_+^{\sum_j K_j}$ as follows.

$$\mathbf{C} = \left\{ \sum_{\mathbf{u}} \Gamma(\mathbf{u}) \mathbf{a}(\mathbf{u}) : \mathbf{a}(\mathbf{u}) \in C^{\text{cvx}}(\mathbf{u}) \right\}$$

Based on this we define another set $\mathcal{C} \subset \mathbb{R}_+^N$. Let for any $\mathbf{a} \in \mathbb{R}_+^N$, $\mathbf{a}^E := ((a_1, a_1, \ldots, K_1 \text{ times}), (a_2, a_2, \ldots, K_2 \text{ times}), \ldots, (a_N, a_N, \ldots, K_N \text{ times}))$. Then $\mathcal{C} = \{\mathbf{a} : \mathbf{a}^E \in \mathbf{C}\}$. This set characterizes the capacity region of the crowd system.

Theorem 1. *Any arrival rate $\boldsymbol{\lambda}$ is stabilizable if for some $\epsilon > 0$, $\boldsymbol{\lambda} + \epsilon \mathbf{1} \in \mathcal{C}$ and no arrival rate $\boldsymbol{\lambda}$ can be stabilized if $\boldsymbol{\lambda}$ is outside the closure of the set $\mathcal{C}$.*

Note that we ignore the precedence constraint in defining $C(\mathbf{u})$. This does not conflict with the fact the capacity region is a subset of $\mathcal{C}$, but it may not be obvious $\mathcal{C}$ is in fact the capacity region. *A fortiori*, we show this with a

scheme that respects precedence constraints and stabilizes any rate in the interior of $\mathcal{C}$.

A. Centralized Allocation

Let us develop a statistics-agnostic scheme that stabilizes any arrival rate $\boldsymbol{\lambda}$.

Let $Q_{j,k}(t)$ be the number of unallocated (j,k) steps just before allocation epoch t. This includes steps not allocated at epoch $t-1$ and steps that became available for allocation between $t-1$ and t. Thus, if for any (j,k), $D_{j,k}(t)$ (j,k)-steps were allocated at epoch t and $A_{j,k}(t+1)$ new (j,k)-steps became available between t and $t+1$,

$$Q_{j,k}(t+1) = Q_{j,k}(t) - D_{j,k}(t) + A_{j,k}(t+1).$$

Note that, for any j and $K_j \geq k > 1$, new (j,k)-steps become available only when some $(j,k-1)$ steps have been completed. Service times $\{r_{j,k,s}\}$ are strictly less than the duration between two allocation epochs. So, any step allocated at epoch t is completed before epoch $t+1$. Hence, for any j and $K_j \geq k > 1$: $A_{j,k}(t+1) = D_{j,k-1}(t)$. On the other hand, for any j and $k = 1$, we have an external arrival $A_j(t+1)$ between epoch t and $t+1$.

At any time t, for a given resource availability, an allocation rule determines resources to be allocated for certain number of (j,k)-steps. We denote this by $S_{j,k}(t)$. Note that $D_{j,k}(t) = \min(Q_{j,k}(t), S_{j,k}(t))$. Our goal is to design a scheme that finds a good $\{S_{j,k}(t)\}$ for a given $\{Q_{j,k}(t)\}$ and $\mathbf{U}(t) = \mathbf{u}$.

Centralized Allocation

Input: $\{Q_{j,s}(t) : j \in [N], s \in [S]\}$ and $\mathbf{U}(t)$ at t

Output: $\{S^*_{j,k}(t)\}$ and allocation of steps to agents

1: Define: $l_{j,r}$: number of leaves in the subtree of T_j rooted at r

2: Obtain $\{S^*_{j,k}(t)\} = \arg\max_{s_{j,k} \in C(\mathbf{U}(t))}$

$$\sum_j \sum_{k=1}^{K_j} \sum_{r:k \to r \in T_j} s_{j,k} l_{j,r} (Q_{j,k}(t) - Q_{j,r}(t)) \tag{4}$$

3: For each (j,k) allocate $S^*_{j,k}(t)$ (j,k)-steps

This allocation scheme is statistics-agnostic and explicit in terms of system state. Also, note that by the design of the scheme the precedence constraint is automatically satisfied. One important thing to note is that the allocation scheme is generic, in the sense that this policy can be easily adapted for different agent- and step-flexibility. Note that $C(\mathbf{U}(t))$ comes from the allocation constraints of the system. If in (4) we replace $C(\mathbf{U}(t))$ by the corresponding allocation set, the centralized algorithm becomes a generic allocation routine.

In fact, the generic statistics-agnostic routine for centralized allocation scheme described above is optimal, in the sense that any arrival rate that can possibly be stabilized by any policy can also be stabilized by this scheme.

Theorem 2. *The centralized allocation routine described above stabilizes any $\boldsymbol{\lambda}$ if $\boldsymbol{\lambda} + \epsilon \mathbf{1} \in \mathcal{C}$, the capacity region of the corresponding system for any $\epsilon > 0$.*

Though the scheme has similarity with back-pressure algorithms [12], [18], [19]; unlike the back-pressure scheme it also uses graph parameter ($l_{j,r}$) in computing the weights. Proof is using a Lyapunov function involving $\{l_{j,r}\}$ and queue-lengths.

Instead of directed rooted tree if the precedence constraint is a directed acyclic graph the same results extend. It would be apparent from proof of Theorem 1 that the converse (outer-bound on capacity) depend on the precedence graph. On the other hand, for any precedence constraint given by a directed acyclic graph, there exists a precedence constraint given by a directed rooted tree such that the tree constraint does not violate the directed acyclic graph constraint. Then, by applying the above centralized algorithm for this directed rooted tree capacity can be achieved.

V. Inflexible Agents and Flexible Steps

Here we characterize the limits of tasks allocation where all steps are flexible and agents are inflexible. Sec. III presented a generic capacity characterization and algorithm; this section investigates computational aspects of the generic algorithm for this particular system and also proposes a simple decentralized scheme that works well under a broad class of assumptions.

Consider $C_{I,F}(\mathbf{u})$, the set of possible allocations with inflexible agents (I) and flexible tasks (F) for availability of agents, $\mathbf{u}$. Recall the allocation scenario in Sec. III to determine a generic $C(\mathbf{u})$: A1–A3 are the same for any system flexibility, but A4 is specific. For an (I, F) system we have the following.

To allocate $a_{j,k,s}$ tasks of different types to a collection of R agents of type $\{m_r : r \in [R]\}$ and available hours $\{h_{m_r,s} : r \in [R]\}$ we must satisfy (1):

$$\sum_{j,k} a_{j,k,s} r_{j,k,s} \leq \sum_{r=1}^{R} h_{m_r,s} \text{ for all } s \in [S]. \tag{5}$$

Note that whenever a step is allocated, all tasks in it must be allocated simultaneously. Hence, we can only allocate $a_{j,k,s}$ tasks with $a_{j,k,s} = a_{j,k,s'}$ $\forall s, s' \in [S]$ when satisfying (5).

Given $C_{I,F}(\mathbf{u})$, the capacity region $\mathcal{C}_{I,F}$ is obtained in the same way $\mathcal{C}$ was obtained from $C(\mathbf{u})$ in Sec. III.

The generic centralized allocation routine can be similarly specialized for (I, F) systems: $C(\mathbf{U}(t))$ in (4) of the routine is replaced by $C_{I,F}(\mathbf{U}(t))$. The centralized scheme is computable since $C_{I,F}(\mathbf{U}(t))$ can be written explicitly in terms of $\mathbf{U}(t)$, $\mathbf{r}_{j,k}$, and $\mathbf{h}_m$, but it cannot always be computed in polynomial time. Since any allocation in $C_{I,F}(\mathbf{u})$ must satisfy constraint (5), optimization problem (4) can be written as:

$$\begin{aligned} &\max_{s_{j,k} \in \mathbb{Z}_+} \sum_{j} \sum_{k=1}^{K_{j-1}} w_{j,k} s_{j,k} \\ &\text{s.t. } \sum_{j,k} s_{j,k} r_{j,k,s} \leq \sum_{m} u_m h_{m,s} \text{ for all } s \in [S]. \end{aligned} \tag{6}$$

Note that the solution to the problem does not change if we replace $w_{j,k}$ by $\max(v_{j,k}, 0)$, as optimal schemes never allocate resources to negative $w_{j,k}$. Thus, we assume $w_{j,k} \geq 0$.

Note that (6) is a multi-dimensional knapsack problem, where the number of available items of a given weight and value are unbounded [20]. This problem is known to be NP-hard and without any fully polynomial-time

approximation scheme (FPTAS). A polynomial-time approximation scheme (PTAS) is known, but the complexity is exponential in dimension. Recently extended linear programming (LP) relaxations have been proposed, but have the same issues (see [21] and references therein).

We aim to find a simple and fast distributed scheme, but first propose the following LP relaxation-based, polynomial-time (in N and M) scheme that gives nearly optimal centralized allocation for a large crowd system (under crowd scaling).

$$\begin{aligned} &\{\hat{S}_{j,k}(t)\} = (6) \text{ with relaxing of } \{s_{j,k} \in \mathbb{R}_+\} \\ &\text{Allocate } \{S^r_{j,k}(t) = \lfloor \hat{S}_{j,k}(t) \rfloor\} \end{aligned} \tag{7}$$

We cannot give performance guarantees for this scheme at each allocation epoch for arbitrary $Q_{j,k}$, but for a sufficiently large crowd system, this scheme stabilizes almost any arrival rate that can be stabilized.

Theorem 3. *Under crowd scaling, for any $\alpha < 1$ there is an N_0 such that for any system with $N \geq N_0$, the LP-based scheme* (7) *stabilizes any arrival rate in $\alpha\mathcal{C} = \{\mathbf{a} : \frac{\mathbf{a}}{\alpha} \in \mathcal{C}\}$.*

A. Decentralized Allocation

In this section we develop a simple decentralized scheme with good performance guarantees. As discussed before, often one of the main reasons for customers to go to a crowd platform is the ability to choose workers themselves. As such, we propose a simple greedy scheme that allows customers the freedom of choice with minimal intervention from platform operators. This also reduces the platform's operational cost.

In greedy allocation, each step competes against others to find an allocation for all of its tasks. Contention can be resolved arbitrarily, e.g., random, pre-ordered, or age-based.

The *Prioritized Greedy* algorithm below performs greedy allocation among all steps across all types of tasks that are in the same order. It starts with steps that are in the beginning of the precedence tree and once these steps find an allocation (or cannot be allocated), only then are steps lower in the corresponding precedence trees allowed to allocate themselves.

Algorithm 1 Prioritized Greedy

Define $D = \max_j$ depth of T_j

1: $\mathcal{S}_j = \emptyset$ for all $j \in [N]$
2: **for** d=1:D **do**
3: $\quad \mathcal{S}_j = \{k_j :$ depth of k_j in $T_j = d\}$
4: $\quad$ Greedy allocation among $\cup_j \{j, k_j : k_j \in \mathcal{S}_j\}$ steps
5: **end for**

This algorithm can be efficiently implemented on a crowdsourcing platform with minimal intervention from the platform operator. The operator need only tag unallocated steps in the system based on their depth in the rooted

precedence tree and only show available workers to them after steps at lower depth have exercised their allocation choice. This may be implemented by personalizing the platform's search results.

The algorithm is fast and has good performance guarantees under certain broadly-used assumptions on arrival and availability processes.

Definition 5. *A random variable X is* Gaussian-dominated *if $\mathbf{E}[X^2] \leq \mathbf{E}[X]^2 + \mathbf{E}[X]$ and for all $\theta \in \mathbb{R}$, $\mathbf{E}[e^{\theta(X-\mathbf{E}[X])}] \leq e^{\frac{1}{2}\left(\left(\mathbf{E}[X^2]-\mathbf{E}[X]^2\right)\theta^2\right.}$, and* Poisson-dominated *if for all $\theta \in \mathbb{R}$, $\mathbf{E}[e^{\theta(X-\mathbf{E}[X])}] \leq e^{\mathbf{E}[X](e^{\theta}-\theta-1)}$.*

These domination definitions, commonly assumed in bandit problems [22], imply that variation around the mean is dominated in a moment generating function sense by that of a Gaussian (Poisson) random variable. Such a property is satisfied by many distributions used to model arrival processes, including in crowdsourcing systems [23].

Theorem 4. *Consider inflexible agents and flexible steps crowdsourcing systems (size N) where for any s, s' $|\sum_m \mu_m h_{m,s} - \sum_m \mu_m h_{m,s'}|$ is sub-poly(N), i.e., $o(N^\delta), \forall \delta > 0$, arrival and availability processes are Poisson-dominated (and/or Gaussian-dominated), and system scales as per crowd-scaling. Then, for any $\alpha \in (0,1)$, $\exists N_\alpha$ s.t. $\forall N \geq N_\alpha$, any arrival rate $\boldsymbol{\lambda} \in \alpha \mathcal{C}_{I,F}$ can be stabilized by Prioritized Greedy, and at the steady state the total number of unallocated steps in the system across all types is $O(\log N)$ w.p. $1 - O\left(\frac{1}{N^2}\right)$.*

This implies Prioritized Greedy can stabilize almost any stabilizable arrival rate while ensuring the number of unallocated tasks scales more slowly than the system size.

VI. Flexible Agents and Flexible Steps

Now consider systems with flexible agents and flexible steps (F, F), and characterize capacity regions. For a given availability of agents $\mathbf{u}$, the set of possible step allocations are $C_{F,F}(\mathbf{u})$. As for $C(\mathbf{u})$ in Sec. III this satisfies A1–A3 in the allocation scenario; A4 for (F, F) systems is as follows.

A certain number of steps $\{a_{j,k}\}$ of each type can be allocated to a set of agents $\{1, 2, \ldots, R\}$ of types $\{m_r : r \in [R]\}$ if there exists a set of R vectors in $\mathbb{R}^S$, $\tilde{h}_{m_r} = (\tilde{h}_{m_r,s} \geq 0 : s \in [S])$, such that:

$$\begin{aligned}
&\tilde{h}_{m_r,s} = 0 \text{ if } s \notin S_{m_r} \text{ for all } s, m_r,\\
&\sum_s \tilde{h}_{m_r,s} \leq \tilde{h}_{m_r} \text{ for all } r \in [R], \text{ and}\\
&\sum_{j,k} a_{j,k} r_{j,k,s} \leq \sum_{r=1}^{R} \tilde{h}_{m_r,s} \text{ for all } s \in [S]. \qquad (8)
\end{aligned}$$

Based on the set of possible allocations $C_{F,F}(\mathbf{u})$, the capacity region $\mathcal{C}_{F,F}$ can be characterized just as in Sec. III.

Similar to Sec. V, if we replace $C(\mathbf{U}(t))$ by $C_{F,F}(\mathbf{U}(t))$ in the centralized allocation routine we obtain an optimal policy for the (F, F) system. It is not hard to see that for the instance where each agent has exactly one skill, the problem is again a multi-dimensional knapsack problem and therefore NP-hard. We develop a computationally-efficient scheme.

If there are R agents of type $\{m_r : r \in R\}$ available, then the centralized allocation problem at time t is to optimize:

$$\max_{(s_{j,k}\in\mathbb{Z}_+),(\tilde{h}_{m_r,s}\in\mathbb{R}_+)} \sum_{j,k} w_{j,k}s_{j,k} \text{ s.t. constraints in (8).} \tag{9}$$

This is a mixed ILP with $\sum_{j=1}^{N} K_j$ integer variables and RS real variables. The complexity of this problem scales with the number of available agents in the system, R. We would like to avoid such a scaling as R may be much larger than M and N in a crowd system. Hence, we pose another optimization problem where the number of variables scales with M and N.

Given $U(t) = \mathbf{u}$,

$$\begin{aligned}
&\max_{(s_{j,k}\in\mathbb{Z}_+),(\alpha_{m,s}\in\mathbb{R}_+)} \sum_{j,k} w_{j,k}s_{j,k} \\
&\text{s.t. } \alpha_{m,s} = 0 \text{ if } s \notin S_m \text{ for all } s, m, \\
&\sum_{s} \alpha_{m,s} \leq 1 \text{ for all } m \in [M], \text{ and} \\
&\sum_{j,k} s_{j,k}r_{j,k,s} \leq \sum_{m=1:M} u_m h_m \alpha_{m,s} \text{ for all } s \in [S].
\end{aligned} \tag{10}$$

Note that this optimization problem yields an allocation satisfying all constraints for flexible agent allocation. This is because $\alpha_{m,s}$ is the fraction of time of an agent of type m that has been given to skill s, which can be positive only when agent of type m has skill s. The last inequality ensures that the skill-hour constraint per skill is satisfied. Hence, this is a feasible allocation procedure.

This is again a mixed ILP, but with $M + N$ variables. Note that this problem is also NP-hard, corresponding to a multi-dimensional knapsack problem if $|S_m| = 1, \text{ for all } m \in [M]$. We design a centralized scheme that allocates steps based on the following LP relaxation. Given $U(t) = \mathbf{u}$,

$$\begin{aligned}
&(\hat{s}_{j,k} : j, k) = \arg \max_{(s_{j,k}\in\mathbb{R}_+),(\alpha_{m,s}\in\mathbb{R}_+)} \sum_{j,k} w_{j,k}s_{j,k} \\
&\text{s.t. constraints in (10) and allocate } \{\lfloor \hat{s}_{j,k} \rfloor\} \text{ steps.}
\end{aligned} \tag{11}$$

This scheme has the following performance guarantee.

Theorem 5. *Under crowd scaling, for any $\alpha < 1$ there is an N_0 s. t. for any system with $N \geq N_0$, the LP-based scheme* (11) *stabilizes any arrival rate in $\alpha\, \mathcal{C}_{F,F} = \{\mathbf{a} : \frac{\mathbf{a}}{\alpha} \in \mathcal{C}_{F,F}\}$.*

Proof of this theorem is based on the equivalence of (9) and (10).

A. Decentralized Allocation

Now we develop a decentralized allocation scheme that requires minimal centralized operation, and gives customers the option to choose from a pool of multiple agents.

Algorithm 2 Prioritized Greedy with Flexibility

Initialize: $\{\gamma(t-t_0) \in [0,1], t \geq t_0\}$, at starting time t_0 $\bar{A}(t_0) = \mathbf{1}$, $\epsilon > 0$

1: Update at each t:

$$\bar{\mathbf{A}}(t) = (1-\gamma(t-t_0))\bar{\mathbf{A}}(t-1) + \gamma(t-t_0)\mathbf{A}(t)$$
$$\bar{\mathbf{U}}(t) = (1-\gamma(t-t_0))\bar{\mathbf{U}}(t-1) + \gamma(t-t_0)\mathbf{U}(t)$$

2: Solve for $\{\psi_{m,s}(t) : s \in [S], m \in [M]\}$

$$\begin{aligned} &\max 1 \text{ s.t.} \\ &\sum_{j,k} \bar{A}_j(t) r_{j,k,s} \leq (1-\epsilon) \sum_{m: s \in S_m} \bar{U}_m(t) \psi_{m,s}(t) \\ &\psi_{m,s}(t) \geq 0, \sum_s \psi_{m,s}(t) \leq 1 \\ &\psi_{m,s}(t) > 0 \text{ only if } s \in S_m, \end{aligned} \tag{12}$$

if no solution pick $\psi_m(t)$ randomly from a simplex in $\mathbb{R}^S$.

3: Initialize sets: $\mathcal{B}_s = \emptyset, \forall s \in [S]$

4: For each type m: put each available agent in one of $\{\mathcal{B}_s\}$ w. p. $\{\psi_{m,s}(t)\}$ (independent rolls of loaded dices)

5: Create *inflexible* agents: an agent of type m in $\mathcal{B}_s$ has h_m available time only for skill s

6: Run **Prioritized Greedy** for this (I,F) system

This algorithm is amenable to crowdsourcing platform implementation. Note $\bar{\mathbf{A}}(t)$ is available from recent history. Creating the set $\mathcal{B}_s$ is simple: for any agent of type m we just randomly tag (as per ψ) with a particular skill and it is shown only tasks with this particular skill. Similarly customers are only shown that the agent has only the particular skill. The rest of the algorithm is exactly like Prioritized Greedy where we create classes of steps and priorities among them and then within each class the allocation is arbitrarily greedy.

We can guarantee Alg. 2 performance when γ satisfies: $\gamma(x) = O\left(x^{-1}\right)$ and $\gamma(x) = \Omega\left(x^{-\frac{1}{2}+\epsilon}\right), \epsilon > 0$.

Theorem 6. *Consider a flexible agents and flexible steps crowdsourcing system with availability processes that are Poisson (and/or Gaussian) dominated with restricted asymmetry, i.e.,* $\max_{s,s'} |\sum_{j,k} \lambda_j r_{j,k,s} - \sum_{j,k} \lambda_j r_{j,k,s'}|$, *being* O *(subpoly*(N)*). For any* $\alpha \in (0,1)$, $\exists N_\alpha$ *s.t.* $\forall N \geq N_\alpha$ *in such systems of size* N *that follow crowd scaling any arrival rate* $\boldsymbol{\lambda} \in \alpha \mathcal{C}_{F,F}$ *can be stabilized by Alg. 2 and at the steady state (i.e., for any finite* t *when* $t_0 = -\infty$*) the total number of unallocated steps in the system across all types is* $O(\log N)$ *w.p.* $1 - O\left(\frac{1}{N^2}\right)$.

VII. Flexible Agents and Inflexible Steps

Now consider the setting where agents may split their available service-time across their skills, but a step must be allocated to one agent. Multiple agents cannot pool their service time to serve a step. As before, for an agent

availability vector $\mathbf{u}$, there is a set of possible allocations of steps (of different (j,k)-types) to agents, denoted $C_{F,I}(\mathbf{u})$. Given $C_{F,I}(\mathbf{u})$ and the distribution of agent availability $\Gamma(\mathbf{u})$, we can define a capacity region $\mathcal{C}_{F,I}$ in the same way as $\mathcal{C}_{I,F}$ is defined in Sec. V based on $C_{I,F}$. Similarly, the generic centralized routine can be adapted by changing the optimization over $C(\mathbf{U}(t))$ to an optimization over $C_{F,I}(\mathbf{U}(t))$ while ensuring optimality of the modified scheme for (F,I) system.

Allocation constraint (2) is for allocation of steps to a particular agent. For a given set of agents of different types $\mathbf{u} = (u_1, u_2, \ldots, u_M)$ the allocation constraint can be written based on (2). Note that for inflexible steps agents cannot pool service-times to serve a step. Consider a set of available agents $a_1, a_2, \ldots, a_R$, of types $1, 2, \ldots, m_r$ respectively. An allocation of $\{s_{j,k} \in \mathbb{Z}_+ : k \in [K_j], j \in [N]\}$ steps to these agents is possible if and only if there are integers $\{z_{j,k,r} \in \mathbb{Z}_+\}$ such that $z_{j,k,r}$ (j,k)-steps are allocated to agent a_r and all $\{s_{j,k}\}$ steps are allocated to some agent, i.e., for each r there is an α_r in an S-dimensional simplex so that:

$$\sum_{j,k} \mathbf{r}_{j,k} z_{j,k,r} \leq \alpha_r h_{m_r}, \alpha_{r,s} = 0 \text{ if } s \notin S_{m_r}, r \in [R]$$

$$\sum_{r} z_{j,k,r} \geq s_{j,k} \forall j,k.$$

Hence, the optimization problem in the centralized allocation routine for (F,I) system is an integer LP of the form:

$$\begin{aligned}
&\max_{s_{j,k} \in \mathbb{Z}_+} \sum_j \sum_{k=1}^{K_j} w_{j,k} s_{j,k} \\
&\text{s.t. } \sum_{j,k} \mathbf{r}_{j,k} z_{j,k,r} \leq \alpha_r h_{m_r}, \alpha_{r,s} = 0 \text{ if } s \notin S_{m_r}, r \in [R], \\
&\quad \sum_r z_{j,k,r} \geq s_{j,k} \text{ for all } j,k, \text{ and} \\
&\quad z_{j,k,r} \in \mathbb{Z}_+ \text{ for all } j,k,r, R = \sum_m u_m.
\end{aligned} \tag{13}$$

Note that like (6), the objective can be written as $\sum_j \sum_{k=1}^{K_{j-1}} v_{j,k} s_{j,k}$ where $v_{j,k} = \max(w_{j,k}, 0)$.

This problem has a special structure which leads to a computationally-efficient algorithm. Consider the following.

$$\begin{aligned}
&\max_{z_{j,k,r} \in \mathbb{Z}_+} \sum_j \sum_{k=1}^{K_j} v_{j,k} \sum_r z_{j,k,r} \\
&\text{s.t. } \sum_{j,k} \mathbf{r}_{j,k} z_{j,k,r} \leq \alpha_r h_{m_r}, \alpha_{r,s} = 0 \text{ if } s \notin S_{m_r}, r \in [R], \text{ and} \\
&\quad z_{j,k,r} \in \mathbb{Z}_+ \text{ for all } j,k,r, R = \sum_m u_m.
\end{aligned} \tag{14}$$

When operating at the optimum of (13), $\sum_r z_{j,k,r} = s_{j,k}$, and so we see that (13) and (14) have the same optimal value. Hence, we solve problem (14) instead of problem (13).

Note that as there is no constraint between $\{z_{j,k,r} : j,k\}$ and $\{z_{j,k,r'} : j,k\}$, problem (14) decomposes into $\sum_m u_m$ optimization problems, each for an available agent. Consider the optimization problem for an agent of type

m.

$$\max_{z_{k,j}} \sum_j \sum_{k=1}^{K_j} v_{j,k} z_{j,k}$$
$$\text{s.t.} \sum_{j,k} \mathbf{r}_{j,k} z_{j,k} \leq \alpha_r h_{m_r}, \alpha_{r,s} = 0 \text{ if } s \notin S_{m_r}, \tag{15}$$

which is again equivalent to the following problem, expressed in terms of the set S_m of skills of type m agent:

$$\max_{z_{k,j}} \sum_{j,k: r_{j,k,s}=0 \text{ if } s \notin S_m} v_{j,k} z_{j,k}$$
$$\text{s.t.} \sum_{j,k: r_{j,k,s}=0 \text{ if } s \notin S_m} z_{j,k} \sum_s r_{j,k,s} \leq h_m. \tag{16}$$

This is a one-dimensional knapsack problem, and there are dynamic programming (DP) pseudo-polynomial time algorithms for solving. Since the sack size h_m is finite (does not scale with the system), the DP has computational complexity $O(\sum_j K_j)$. This implies the centralized scheme decomposes into $\sum_m u_m$ problems, each of which can be solved in polynomial time.

Thus, the centralized scheme naturally leads to a decentralized scheme where each available agent solves (16) and uses the optimal solution as its *potential* allocated steps. Agents may use an arbitrary contention mechanism among themselves to decide which agent allocates first. Upon resolving contention, agents pick steps greedily by solving (16). Since the decentralized scheme follows directly from the centralized one (13), performance guarantees from Thm. 2 hold.

Although this simple decentralized scheme is optimal, it does not give customers freedom of choosing agents. Thus, we propose another decentralized scheme where customers get to pick any agent from a subset of available agents.

Alg. 3 allows the different types of steps to pick agents greedily, but in a restricted manner. It prioritizes steps with lower depth like Prioritized Greedy. Among steps with the same priority (in terms of depth), it gives preference to steps requiring more skills to ensure an agent with multiple skills is not used unwisely for a step with lesser requirements.

Theorem 7. *Consider a flexible agent and inflexible steps crowdsourcing system where each type of agent has $O(1)$ skills, and arrival as well as availability processes are Poisson (and/or Gaussian) dominated, $\{S_m : m \in [M]\}$ is a partition of $[S]$ and $\sum_s r_{j,k,s}$ are same for all (j,k). For this system for any $\alpha \in (0,1)$, $\exists N_\alpha$ s.t. $\forall N \geq N_\alpha$ in such systems of size N that follows crowd scaling any arrival rate $\boldsymbol{\lambda} \in \alpha \mathcal{C}_{F,I}$ can be stabilized by Alg. 3 and at steady-state the total number of unallocated steps in the system across all types is $O(\log N)$ w.p. $1 - O\left(\frac{1}{N^2}\right)$.*

For many systems the total sizes of steps are nearly identical and so the assumption on total size is not restrictive, though results can be extended to the case where the total sizes are random with the same mean. The assumption $\{S_m : m \in [M]\}$ is a partition is required for proving the performance guarantee, but the algorithm (actually a simpler version) works well on simulations. The above performance guarantee can be extended for the following

Algorithm 3 Restricted Greedy

Compute and Store: one-time

for $d = 1 : D$ **do**
  $\mathcal{A}_d = \cup_j \{$set of required skills for steps at depth d of $T_j\}$
  $p_d = 0$
  while $\mathcal{A}_d \neq \emptyset$ **do**
    $p_d = p_d + 1$
    Pick maximal subsets from the collection $\mathcal{A}_d$, say $\mathcal{L}^d_{p_d}$.
    $\mathcal{A}_d = \mathcal{A}_d \backslash \mathcal{L}^d_{p_d}$
  end while
end for

Allocation at time t

Define $D = \max_j$ depth of T_j

for $d = 1 : D$ **do**
  while $k \leq p_d$ **do**
    Steps in $\mathcal{L}^d_k$ allocate themselves greedily (ties are broken arbitrarily)
    $k = k + 1$
  end while
end for

conditions. $\{S_i : i \in \mathcal{I}\}$ is a partition of $[S]$ for some $\mathcal{I} \subset [M]$ and for any m $S_m \subset S_i$ for some $i \in \mathcal{I}$, $D = 1$, and for any (j, k), (j', k') pair, $\{s : r_{j,k,s} > 0\}$ and $\{s : r_{j',k',s} > 0\}$ either have no intersection or one is a subset of the other.

VIII. Evaluation

Secs. III–VII characterized limits of different types of crowdsourcing systems, proposed efficient policies for optimal centralized allocation and designed decentralized schemes with provable bounds on backlog while giving customers freedom of choice. This section complements theoretical results by studying real data from Samasource, a non-profit crowdsourcing company and realistic Monte Carlo simulations. We study performance of simplified (in implementation and computation) versions of proposed decentralized algorithms above.

Let us first describe evaluating allocation using real data. The dataset contains 9.3M tasks and each belongs to a specific project. Some projects are regarded as *real-time* which means they have higher priority. The overall number of tasks that belong to the real-time projects is about 4.2M. Each task comprises 1 or 2 steps which in turn comprises a single substep. Some tasks have strict step ordering, i.e., the previous step must be completed before the next could be scheduled. Average substep working time requirement is 340 sec. From the data, we calculate

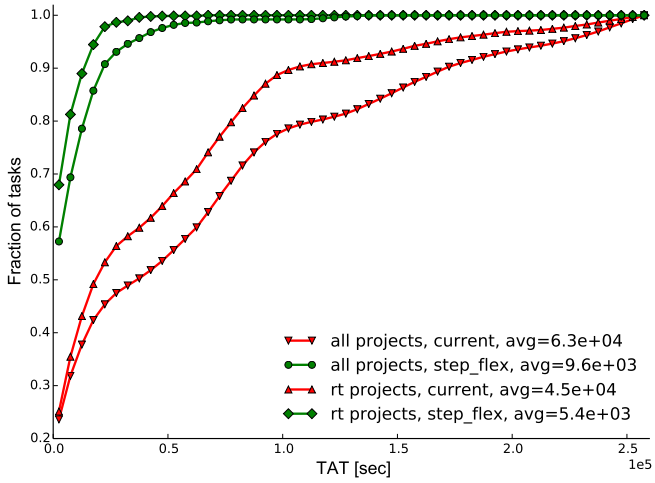

Fig. 1. CDF of tasks turn-around time (TAT) using real dataset. Current allocation on the platform "current" vs our algorithm "STEP_FLEX".

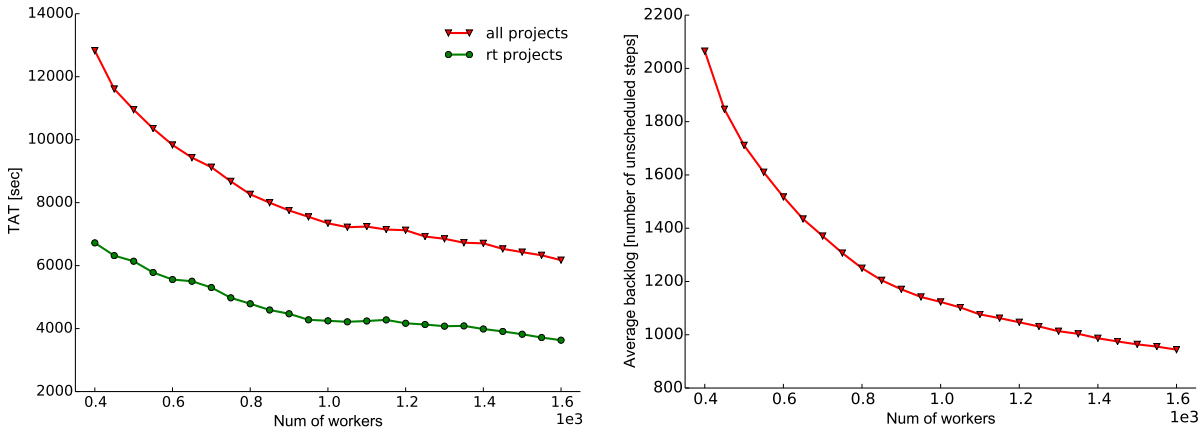

(a) (b)

Fig. 2. Performance of our STEP_FLEX algorithm on real data, as a function of number of workers. (a) Tasks turn-around time (TAT). (b) Average backlog (number of unallocated steps in the system).

the turn-around time (TAT) for each task, i.e., the time since the task arrived to the system until the time its last step was completed. The cdfs of TAT for all projects and for real-time projects only are given in Fig. 1.

SamaHub, the platform of Samasource considers both agents and steps to be flexible. We implement a simplified version of the relevant decentralized algorithm, called STEP_FLEX, where we prioritize the steps with higher precedence to choose agents greedily with random tie-breaking.

To compare current allocation on SamaHub with our approach, we use real data as input to STEP_FLEX. Since we lack exact knowledge of worker availability, we make the following assumption in consultation with Samasource.

The number of active workers in the system is 625, evenly distributed across four time zones: $-4, 0, 3, 5.5$, where each worker works every day from 9am to 5pm. Each worker possesses the skills required for any substep in the dataset. Fig. 1 compares the cdf of TAT of our approach STEP_FLEX (simulated with the data as input) with currently deployed scheme. Our algorithm substantially outperforms current scheme: average TAT for all projects is $\times 6.5$ better and more than $\times 8$ better for real-time projects. This improvement is also influenced by our implementation, which is not restricted by the currently-practiced organizational structure.

Fig. 2 shows how STEP_FLEX performs as a function of number of workers. As the number of workers grows, TAT decreases (see Fig. 2(a)). The benefit of adding more workers can be seen even more clearly when analyzing backlog, i.e., the average number of steps that entered the system but not yet scheduled, see Fig. 2(b).

We also evaluate our algorithms on synthetic data, considering flexible agents and flexible steps, and flexible agents and inflexible steps. Algorithm STEP_FLEX is used for the first system and a simplified version of the Restricted Greedy scheme, STEP_INFLEX, where we prioritize steps with higher skill requirements and allocate among them greedily is used for the second. We also consider a scenario in between flexible and inflexible steps, where each substep is allocated to a single agent, but different substeps of a step can be allocated to different agents. For this, we develop STEP_SEMIFLEX where steps allocate themselves greedily while ensuring a substep gets all service from an agent. We expected STEP_FLEX to outperform STEP_INFLEX, but we found somewhat surprisingly that STEP_FLEX and STEP_SEMIFLEX perform very similarly.

The first set of generated data has tasks with up to three steps in each and with strict ordering. Each step comprises one to three random substeps out of five possible types. Working time requirement for each substep is uniformly distributed between 60 and 600 sec. Each worker in the system has daily availability from 9am to 5pm, evenly distributed across four time zones: $-4, 0, 3, 5.5$. A worker possesses a random set of skills that enables her to work on up to three (out of five) substep types. For each of our three algorithms we compare three metrics: TAT, backlog queue, and worker utilization. The experiment simulated a single run over a timespan of 40 days.

Fig. 3 shows algorithms STEP_FLEX and STEP_SEMIFLEX outperform STEP_INFLEX for both cases: 500 workers in the system and 700 workers. When the load on the system is 150 tasks/hour and the number of workers is 500, algorithm STEP_INFLEX is substantially worse since it becomes unstable for this load. Notice that STEP_FLEX and STEP_SEMIFLEX perform very similarly, which can be explained by relatively short substep work time requirement (in which case splitting becomes a rare event). Also note that worker utilization of STEP_INFLEX is not much worse than of the other algorithms. This can be explained by the long backlog queue of STEP_INFLEX. Though it is harder for STEP_INFLEX to find a worker capable of working on the whole step, when the backlog becomes large, the probability that a given worker will be assigned to some whole step grows.

The last set of results uses the same synthetic data as before, but the working time requirement for each substep is now uniformly distributed between 600 and 6000 sec. Fig. 4 shows a slight advantage of STEP_FLEX over STEP_SEMIFLEX. Due to the longer working time requirements per substep, cases in which a substep may be split to improve allocation are more probable. In this scenario, the disadvantage of STEP_INFLEX is more obvious: for a load of 50 tasks/hour and 1200 workers, its TAT and backlog are very large and unstable.

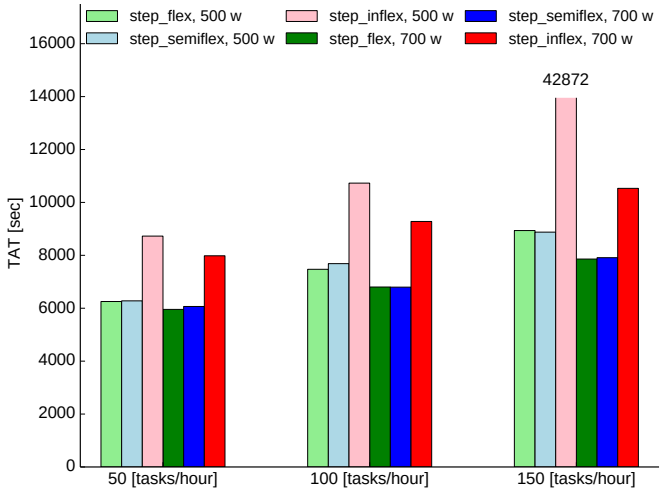

(a)

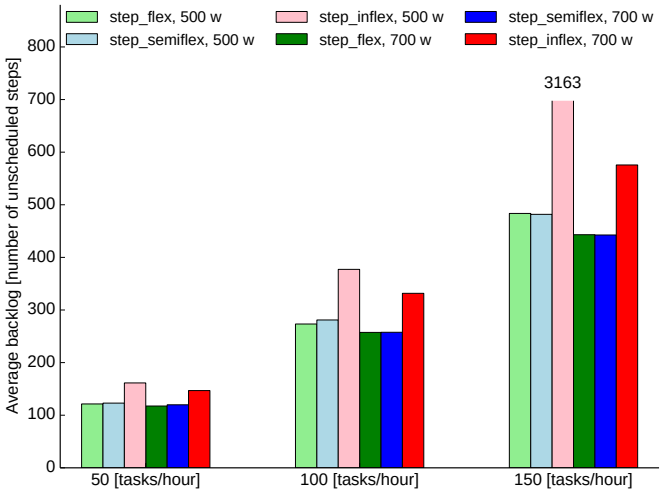

(b)

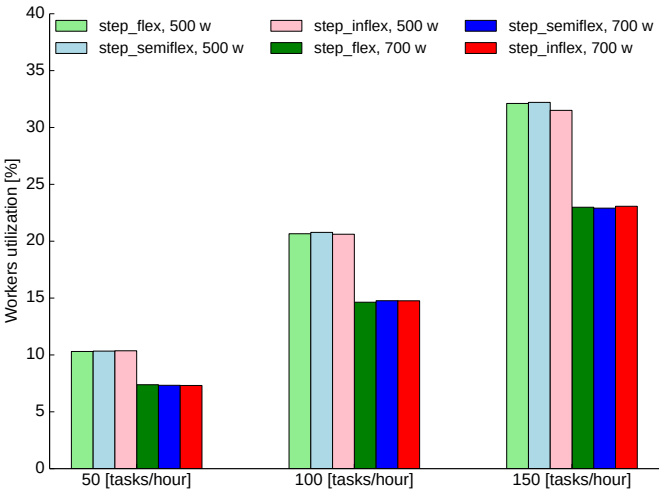

(c)

Fig. 3. Performance of our algorithms on synthetic data with *short* sub-steps ($60 - 600$ sec), as a function of load. (a) Tasks turn-around time (TAT). (b) Average backlog (number of unscheduled steps in the system). (c) Workers utilization.

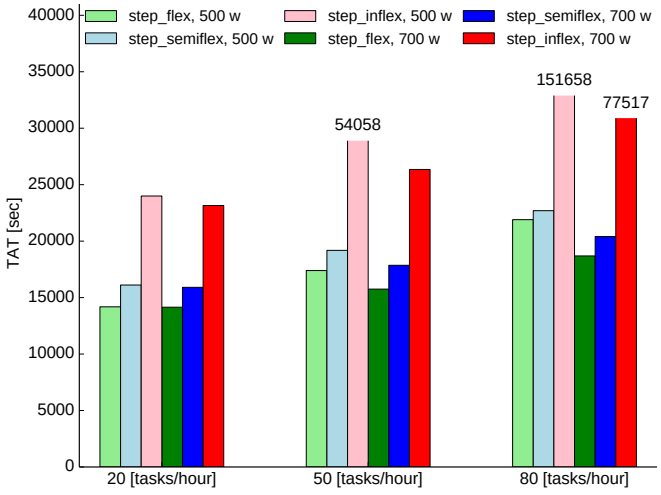

(a)

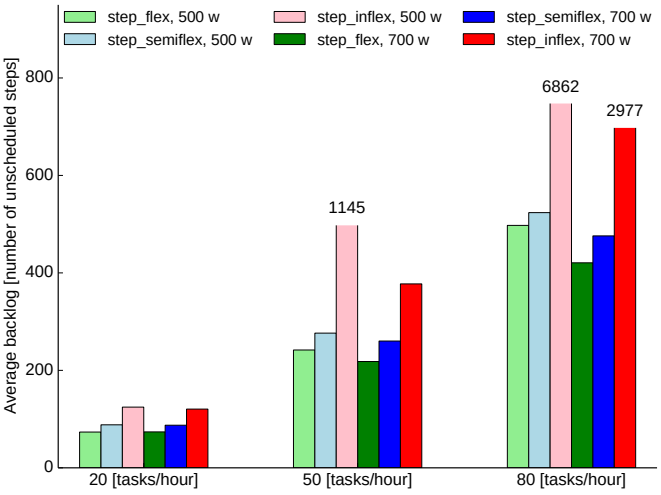

(b)

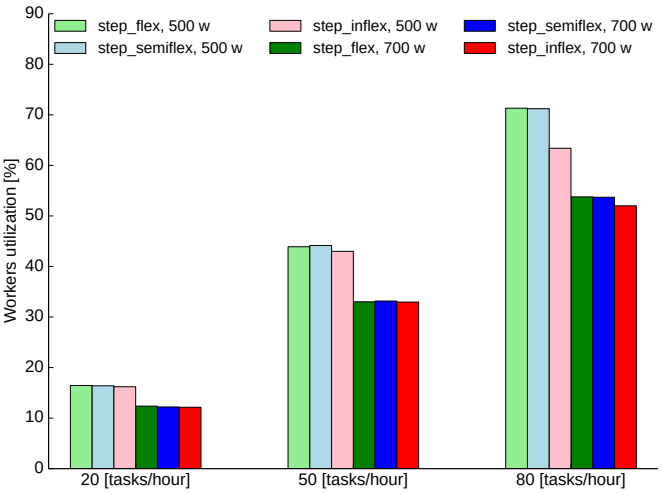

(c)

Fig. 4. Performance of our algorithms on synthetic data with *long* sub-steps ($600 - 6000$ sec), as a function of load. (a) Tasks turn-around time (TAT). (b) Average backlog (number of unscheduled steps in the system). (c) Workers utilization.

To summarize, our approach substantially outperforms Samasource's current allocation scheme. While STEP_FLEX achieves best performance in terms of TAT and backlog, STEP_SEMIFLEX may be a good alternative. Its performance is almost the same but does not require splitting substeps among different workers, and is computationally lighter.

IX. Conclusion

Inspired by skilled crowdsourcing systems, we have developed new algorithms for allocating tasks to agents while handling novel system properties such as vector-valued service requirements, precedence and flexibility constraints, random and time-varying resource availability, large system size, need for simple decentralized schemes requiring minimal actions from the platform provider, and the freedom of customers to choose agents without compromising system performance. We have provided capacity regions, asymptotic performance guarantees for decentralized algorithms, and demonstration of efficacy in practical regimes, via large-scale data from a non-profit crowdsourcing company.

References

[1] A. Chatterjee, M. Borokhovich, L. R. Varshney, and S. Vishwananth, "Efficient and flexible crowdsourcing of specialized tasks with precedence constraints," in *Proc. 2016 IEEE INFOCOM*, Apr. 2016 (to appear).

[2] A. Cuenin, "Each of the top 25 best global brands has used crowdsourcing," Jun. 2015. [Online]. Available: http://www.crowdsourcing.org/editorial/each-of-the-top-25-best-global-brands-has-used-crowdsourcing/50145

[3] D. Tapscott and A. D. Williams, *Wikinomics: How Mass Collaboration Changes Everything*, expanded ed. New York: Portfolio Penguin, 2006.

[4] F. Gino and B. R. Staats, "The microwork solution," *Harvard Bus. Rev.*, vol. 90, no. 12, pp. 92–96, Dec. 2012.

[5] A. Marcus and A. Parameswaran, "Crowdsourced data management: Industry and academic perspectives," *Foundations and Trends in Databases*, vol. 6, no. 1-2, pp. 1–161, Dec. 2015.

[6] T. W. Malone, R. Laubacher, and C. Dellarocas, "The collective intelligence genome," *MIT Sloan Manage. Rev.*, vol. 51, no. 3, pp. 21–31, Spring 2010.

[7] K. J. Boudreau and K. R. Lakhani, "Using the crowd as an innovation partner," *Harvard Bus. Rev.*, vol. 91, no. 4, pp. 60–69, Apr. 2013.

[8] S. Dustdar and M. Gaedke, "The social routing principle," *IEEE Internet Comput.*, vol. 15, no. 4, pp. 80–83, July-Aug. 2011.

[9] D. DiPalantino and M. Vojnović, "Crowdsourcing and all-pay auctions," in *Proc. 10th ACM Conf. Electron. Commer. (EC'09)*, Jul. 2009, pp. 119–128.

[10] L. R. Varshney, S. Agarwal, Y.-M. Chee, R. R. Sindhgatta, D. V. Oppenheim, J. Lee, and K. Ratakonda, "Cognitive coordination of global service delivery," arXiv:1406.0215v1 [cs.OH]., Jun. 2014.

[11] J. Kleinberg and É. Tardos, *Algorithm Design*. Addison-Wesley, 2005.

[12] R. Srikant and L. Ying, *Communication Networks: An Optimization, Control and Stochastic Networks Perspective*. Cambridge University Press, 2014.

[13] M. L. Pinedo, *Scheduling: Theory, Algorithms, and Systems*. Springer, 2012.

[14] G. Pang and A. L. Stolyar, "A service system with on-demand agent invitations," *Queueing Systems*, Nov. 2015.

[15] A. Chatterjee, L. R. Varshney, and S. Vishwananth, "Work capacity of freelance markets: Fundamental limits and decentralized schemes," in *Proc. 2015 IEEE INFOCOM*, Apr. 2015, pp. 1769–1777.

[16] F. A. Chudak and D. B. Shmoys, "Approximation algorithms for precedence-constrained scheduling problems on parallel machines that run at different speeds," *J. Algorithms*, vol. 30, no. 2, pp. 323–343, Feb. 1999.

[17] R. Pedarsani, "Robust scheduling for queueing networks," Ph.D. dissertation, University of California, Berkeley, Berkeley, CA, 2015.

[18] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.

[19] M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems*. Morgan & Claypool, 2010.

[20] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*. Springer, 2004.

[21] D. A. G. Pritchard, "Linear programming tools and approximation algorithms for combinatorial optimization," Ph.D. dissertation, University of Waterloo, 2009.

[22] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Found. Trends Mach. Learn.*, vol. 5, no. 1, pp. 1–122, Dec. 2012.

[23] M. Vukovic and O. Stewart, "Collective intelligence applications in IT services business," in *Proc. IEEE 9th Int. Conf. Services Comput. (SCC)*, Jun. 2012, pp. 486–493.

[24] A. Chatterjee, L. R. Varshney, and S. Vishwananth, "Work capacity of freelance markets: Fundamental limits and decentralized schemes," arXiv:1508.00023 [cs.MA], Jul. 2015.

[25] V. S. Borkar, *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, 2008.

APPENDIX A
PROOFS

In this section we present proofs of the main results.

A. Proof of Theorem 1

Here we only prove that any $\boldsymbol{\lambda}$ outside the closure of $\mathcal{C}$ cannot be stabilized by any policy. To prove achievability, it is sufficient to show that there exists a policy that stabilizes any $\boldsymbol{\lambda}$ in the interior of $\mathcal{C}$. Hence, it is sufficient to prove Thm. 2, which we do later.

This proof consists of the following steps. We first compare two systems, the original system in question and another in which there is no precedence constraint among different steps of a job. We claim that on any sample path under any policy for the first system, there exists a policy in the second system so that the total number of incomplete jobs across all job types in the second is a lower bound (sample path-wise) for that in the first. Then we show that the second system cannot be stabilized for a $\boldsymbol{\lambda}$ outside the closure of $\mathcal{C}$ and so the result follows for the first system.

Note that the claim regarding the number of incomplete jobs across all types in the second system being a lower bound on the first system follows by considering the *same* policy for the second system as for the first system.

To proceed, consider the second system, for which we denote the number of unallocated steps of type (j,k) at epoch t by $\hat{Q}_{j,k}(t)$. Now consider the set $\mathcal{C}$. We claim that this set is coordinate convex, i.e., it is a convex set and if $\mathbf{a}+\epsilon \in \mathcal{C}$ for some $\epsilon, \mathbf{a} \in \mathbb{R}_+^N$ then $\mathbf{a} \in \mathcal{C}$. To prove this claim, we first show that the set $\mathbf{C}$ is coordinate convex.

First we prove that $\mathbf{C}$ is a convex subset of $\mathbb{R}_+^N$. If $\mathbf{a}, \mathbf{a}' \in \mathbf{C}$, then there exist $\left(\mathbf{a}(\mathbf{u}) \in C^{\text{cvx}}(\mathbf{u}) : \mathbf{u} \in \mathbb{Z}_+^M\right)$ and $\left(\mathbf{a}'(\mathbf{u}) \in C^{\text{cvx}}(\mathbf{u}) : \mathbf{u} \in \mathbb{Z}_+^M\right)$ such that

$$\sum_{\mathbf{u}} \Gamma(\mathbf{u})\mathbf{a}(\mathbf{u}) = \boldsymbol{\lambda}, \quad \sum_{\mathbf{u}} \Gamma(\mathbf{u})\mathbf{a}'(\mathbf{u}) = \boldsymbol{\lambda}'.$$

Thus for any $\gamma \in [0,1]$,

$$\gamma\mathbf{a} + (1-\gamma)\mathbf{a}' = \sum_{\mathbf{u}} \Gamma(\mathbf{u})(\gamma\mathbf{a}(\mathbf{u}) + (1-\gamma)\mathbf{a}'(\mathbf{u}).$$

Note that $C^{\text{cvx}}(\mathbf{u})$ is convex since it is the convex hull of $C(\mathbf{u})$; hence $\gamma\mathbf{a}(\mathbf{u}) + (1-\gamma)\mathbf{a}'(\mathbf{u}) \in \mathcal{C}(bu)$, which in turn implies $\gamma\mathbf{a} + (1-\gamma)\mathbf{a}' \in \mathbf{C}$.

For coordinate convexity note that any $\mathbf{a}$ is a $\Gamma(\mathbf{u})$ combination of some $\{\mathbf{a}(\mathbf{u}) \in C^{\text{cvx}}(\mathbf{u})\}$ and any $\mathbf{a}(\mathbf{u})$ is some convex combination of elements of $C(\mathbf{u})$. Also, from the allocation constraints it is apparent that if $\mathbf{a} \in C(\mathbf{u})$ then also $\mathbf{a}' \in C(\mathbf{u})$ if $\mathbf{a}' \leq \mathbf{a}$. These two imply that for any $\mathbf{a} \in \bar{\mathcal{C}}$, if there exists an $\mathbf{a}' \leq \mathbf{a}$ (component-wise) and $\mathbf{a}' \geq \mathbf{0}$, then $\mathbf{a}' \in \bar{\mathcal{C}}$. Hence, $\mathbf{C}$ is coordinate convex.

Note that $\mathcal{C} = \{\mathbf{a} : \mathbf{a}^E \in \mathbf{C}\}$. Note that if $\mathbf{a} \leq \mathbf{a}'$ (coordinate-wise) then the same is true for $\mathbf{a}^E$ and $\mathbf{a}'^E$. Also, if $\mathbf{a}'' = \gamma\mathbf{a} + (1-\gamma)\mathbf{a}'$ for any $\gamma \in [0,1]$, then $\mathbf{a}''^E = \gamma\mathbf{a}^E + (1-\gamma)\mathbf{a}'^E$. This proves that $\mathcal{C}$ is coordinate convex subset of $\mathbb{R}_+^N$.

As $\mathcal{C}$ is coordinate convex, for any $\boldsymbol{\lambda}$ outside the closure of $\mathcal{C}$ there exists $h \in \mathbb{R}_+^N$ s.t. $h^T\boldsymbol{\lambda} > \sup_{\mathbf{x}\in\mathcal{C}} h^T\mathbf{x}$.

Consider the following. Let $\hat{\mathbf{Q}} = (\hat{Q}_{j,k})$, $\mathbf{A} = (A_{j,k}(t))$ and $\hat{\mathbf{D}}(t) = (D_{j,k}(t))$. Note that as jobs in this system do not have precedence constraints, for all k, $A_{j,k} = A_j(t)$.

$$\begin{aligned}\mathbf{E}\left[h^T\hat{\mathbf{Q}}(t+1)\right] &= \mathbf{E}\left[h^T\left(\hat{\mathbf{Q}}(t) + \mathbf{A}(t) - \hat{\mathbf{D}}(t)\right)\right] \\ &= \mathbf{E}\left[h^T\left|\hat{\mathbf{Q}}(t) + \mathbf{A}(t) - \boldsymbol{\Delta}(t)\right|^+\right]\end{aligned}$$

where $\boldsymbol{\Delta}(t)$ is the number of possible departures under the scheme if there were infinite number of steps of each type, and $|\cdot|^+$ is shorthand for $\max(\cdot, 0)$. As $|x|^+$ is a convex function of x, $h^T\left|\hat{\mathbf{Q}}(t) + \mathbf{A}(t) - \boldsymbol{\Delta}(t)\right|^+$ is a convex function of $\mathbf{Q}(t), \mathbf{A}(t)$, and $\boldsymbol{\Delta}(t)$. Thus by Jensen's inequality:

$$\begin{aligned}&\mathbf{E}\left[h^T\left|\hat{\mathbf{Q}}(t) + \mathbf{A}(t) - \boldsymbol{\Delta}(t)\right|^+\right] \\ &\geq h^T\left|\mathbf{E}\left[\hat{\mathbf{Q}}(t)\right] + \mathbf{E}\left[\mathbf{A}(t)\right] - \mathbf{E}\left[\boldsymbol{\Delta}(t)\right]\right|^+ \\ &\geq h^T\mathbf{E}\left[\hat{\mathbf{Q}}(t)\right] + h^T\mathbf{E}\left[\mathbf{A}(t)\right] - h^T\mathbf{E}\left[\boldsymbol{\Delta}(t)\right].\end{aligned}$$

Note that $h^T\mathbf{E}\left[\boldsymbol{\Delta}(t)\right] \leq \sup_{\mathbf{x}\in\mathcal{C}} h^T\mathbf{x}$ and $h^T\mathbf{E}\left[\mathbf{A}(t)\right] = h^T\boldsymbol{\lambda}$, hence,

$$\mathbf{E}\left[h^T\hat{\mathbf{Q}}(t+1)\right] \geq h^T\mathbf{E}\left[\hat{\mathbf{Q}}(t)\right] + \epsilon,\ \epsilon > 0. \tag{17}$$

So, we have $\mathbf{E}[h^T\hat{\mathbf{Q}}]$ to be unbounded (i.e., for any constant B, there exists a t s.t. $\mathbf{E}[h^T\hat{\mathbf{Q}}(t)] > B$) under any policy. As $h \geq 0$, we have that $\hat{\mathbf{Q}}^T\mathbf{1}$ to be unbounded and hence, the system is not stable under any policy.

B. Proof of Theorem 2

Before proceeding, we state small lemma that will be useful.

Lemma 1. *For any $x, y, z \geq 0$, $(|x-y|^+ + z)^2 \leq x^2 + y^2 + z^2 + 2x(z-y)$.*

Proof:

$$\begin{aligned}(|x-y|^+ + z)^2 &= (|x-y|^+)^2 + z^2 + 2z|x-y|^+ \\ &\leq (|x-y|^+)^2 + z^2 + 2xz \\ &\leq (x-y)^2 + z^2 + 2xz, \\ &= x^2 + y^2 + z^2 + 2x(z-y),\end{aligned}$$

where the last inequality follows because $(\max(0,a))^2 \leq a^2$. ■

Now for the proof of Thm. 2

The process $\{Q_{j,k}(t)\}$ is a discrete-time Markov chain on $\mathbb{Z}_+^{\sum_j K_j}$ under the centralized scheme. This is because arrival and availability processes are i.i.d. and the centralized allocation at epoch t does not depend on process values before t. We show that for this chain, all closed classes are positive recurrent and that the chain enters one of the closed classes almost surely. Note that this implies that starting with any initial distribution, the Markov

chain reaches a stationary distribution (which may depend on the initial condition). This is in the sense that there exists a $d \in \{1, 2, \dots\}$ (as there may be a closed class which is not aperiodic), and a distribution π on $\mathbb{Z}_+^{\sum_j K_j}$ such that $\mathbf{Q}(td) \to \pi$ in distribution.

To show stability we need $\limsup_{t\to\infty} \mathbf{E}[Q_{j,k}(t)] < \infty$, for all (j,k). Towards this, note it is sufficient to show $\mathbf{E}_\pi[\sum_{j,k} Q_{j,k}]$ is finite, because this implies $\lim_{t\to\infty} \mathbf{E}[\sum_{j,k} Q_{j,k}(td)]$ is finite. Note that for any $1 < \tau < d$:

$$\sum_{j,k} Q_{j,k}(td+\tau) \le \sum_{j,k} [Q_{j,k}(t) + \sum_{t'=1}^{d} A_j(t')].$$

Since arrivals have finite expectation, $\limsup_{t\to\infty} \mathbf{E}[Q_{j,k}(t)] < \infty$, for all (j,k).

Now, it is sufficient to prove that starting with any initial distribution, there exists a $d \in \{1, 2, \cdots\}$ such that $\mathbf{Q}(td) \to \pi$ in distribution and $\mathbf{E}_\pi[\sum_{j,k} Q_{j,k}]$ is finite. To prove the convergence in distribution we use a variation of the Foster-Lyapunov theorem presented in [18].

When T_j is a directed rooted tree, we have to consider a Lyapunov function:

$$L(\mathbf{Q}) = \sum_j \sum_k l_{j,k} Q_{j,k}^2,$$

where $l_{j,k}$ is the number of leaves in the subtree of T_j rooted at k.

Before proceeding with the proof of this case, we prove a reordering lemma for the Lyapunov function.

Lemma 2. *For any allocation* $\{S_{j,k}\}$,

$$\sum_j \left(l_{j,1} Q_{j,1}(A_{j,1} - S_{j,1}) + \sum_{k>1} l_{j,k} \left(Q_{j,k}(t) S^*_{j,p_j(k)}(t) - Q_{j,k}(t) S_{j,k}(t) \right) \right)$$
$$= \sum_j \sum_{k=1}^{K_j} \sum_{r \in c_j(k)} l_{j,r} (Q_{j,k}(t) - Q_{j,r}(t))(A_{j,1} - S_{j,k}). \tag{18}$$

Proof: First we claim that for any j,

$$\sum_{k>1} l_{j,k} \left(Q_{j,k}(t) S^*_{j,p_j(k)}(t) - Q_{j,k}(t) S_{j,k}(t) \right) = \sum_{k=2}^{K_j} \sum_{r \in c_j(k)} l_{j,r} (Q_{j,k}(t) - Q_{j,r}(t)) S_{j,k} \tag{19}$$

This can be seen by comparing coefficients of $Q_{j,k}, k > 1$ on both sides of the expression. Note that $\sum_{r \in c_j(k)} l_{j,r} = l_{j,k}$. Also, note that in the sum in the right side, $Q_{j,k}$ appears twice, once in the sum $-\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t)) S_{j,k}$ where the coefficient is $-S_{j,k}$ and again in the sum $-\sum_{r \in c_j(p_j(k))} l_{j,r}(Q_{j,p_j(k)}(t) - Q_{j,r}(t)) S_{j,p_j(k)}$ where the coefficient is $S_{j,p_j(k)}$.

This implies that for any j,

$$- l_{j,1} Q_{j,1} S_{j,1} + \sum_{k>1} l_{j,k} \left(Q_{j,k}(t) S^*_{j,p_j(k)}(t) - Q_{j,k}(t) S_{j,k}(t) \right)$$
$$= - \sum_{k=1}^{K_j} \sum_{r \in c_j(k)} l_{j,r} (Q_{j,k}(t) - Q_{j,r}(t)) S_{j,k}. \tag{20}$$

Note that since $l_{j,k} = \sum_{r \in c_j(k)} l_{j,r}$,

$$l_{j,1}Q_{j,1} = \sum_{r \in c_j(1)} l_{j,r}(Q_{j,1} - Q_{j,r}) + \sum_{r \in c_j(1)} l_{j,r}Q_{j,r}.$$

Again applying the same restructuring of the terms for the subtrees rooted at r, we eventually obtain:

$$l_{j,1}Q_{j,1} = \sum_{r \in c_j(k)} l_{j,r}(Q_{j,k} - Q_{j,r}).$$

The result follows by combining this (multiplied by $A_{j,1}$) with (20). ■

Now for the main proof.

$$\mathbf{E}[L(\mathbf{Q}(t+1))|\mathbf{Q}(t)] = \mathbf{E}\left[\sum_{j,k}(Q_{j,k}(t) - D_{j,k}(t) + A_{j,k}(t))^2|\mathbf{Q}(t)\right].$$

Let $p_j(k)$ denotes the parent of k in T_j and $c_j(k)$ denote the set of children of k in T_j. For all t and $k = 1$, $A_{j,1}(t) = A_j(t)$. For $k > 1$, $A_{j,k}(t) = D_{j,p_j(k)}(t)$. Note that $S^*_{j,k} \geq D_{j,k}$ for all t. So,

$$\begin{aligned}
&\mathbf{E}\left[\sum_{j,k} l_{j,k}(Q_{j,k}(t) - D_{j,k}(t) + A_{j,k}(t))^2|\mathbf{Q}(t)\right] \\
&= \mathbf{E}\left[\sum_{j,k} l_{j,k}(|Q_{j,k}(t) - S^*_{j,k}(t)|^+ + A_{j,k}(t))^2|\mathbf{Q}(t)\right] \\
&\leq \mathbf{E}\left[\sum_j \left(l_{j,1}(|Q_{j,1}(t) - S^*_{j,1}(t)|^+ + A_{j,1}(t))^2 + \sum_{k>1} l_{j,k}(|Q_{j,k}(t) - S^*_{j,k}(t)|^+ + S^*_{j,p_j(k)}(t))^2\right)|\mathbf{Q}(t)\right].
\end{aligned}$$

By Lem. 1,

$$\begin{aligned}
&\mathbf{E}[L(\mathbf{Q}(t+1))|\mathbf{Q}(t)] \\
&= \mathbf{E}\Bigg[\sum_j l_{j,1}\left(A^2_{j,1}(t) + (S^*_{j,1}(t))^2 + Q^2_{j,1} + 2Q_{j,1}(A_{j,1} - S^*_{j,1})\right. \\
&\qquad \left. + \sum_{k>1} l_{j,k}((S^*_{j,k}(t))^2 + (S^*_{j,p_j(k)}(t))^2 + Q^2_{j,k} + 2Q_{j,k}(S^*_{j,p_j(k)} - S^*_{j,k}))\right)|\mathbf{Q}(t)\Bigg] \\
&\leq C_1 + 2\mathbf{E}\left[\sum_{j,k} l_{j,1}(S^*_{j,k}(t))^2|\mathbf{Q}(t)\right] \\
&\quad + 2\mathbf{E}\left[\sum_j \left(l_{j,1}Q_{j,1}(A_{j,1} - S^*_{j,1}) + \sum_{k>1} l_{j,k}Q_{j,k}(S^*_{j,p_j(k)} - S^*_{j,k})\right)|\mathbf{Q}(t)\right] \qquad (21) \\
&\leq C_2 + 2\mathbf{E}\left[\sum_j \left(l_{j,1}Q_{j,1}(A_{j,1} - S^*_{j,1}) + \sum_{k>1} l_{j,k}Q_{j,k}(S^*_{j,p_j(k)} - S^*_{j,k})\right)|\mathbf{Q}(t)\right] \qquad (22)
\end{aligned}$$

Eq. (21) follows because arrival processes have bounded second moments and are i.i.d. and the fact $(S^*_{j,k}(t))^2 \geq 0$

(so, over-counting them gives an upper bound). Eq. (22) is due to the following (where $K = \max_j K_j$):

$$\begin{aligned}
\mathbf{E}\left[\sum_{j,k} l_{j,k}(S^*_{j,k}(t))^2|\mathbf{Q}(t)\right] &\leq K\mathbf{E}\left[\left(\sum_{j,k} S^*_{j,k}(t)\right)^2|\mathbf{Q}(t)\right] \\
&\leq K\frac{1}{\max_{j,k}(\sum_s r_{j,k,s})^2}\mathbf{E}\left[\left(\sum_{j,k,s} S^*_{j,k}(t)r_{j,k,s}\right)^2|\mathbf{Q}(t)\right] \\
&\leq K\frac{1}{\max_{j,k}(\sum_s r_{j,k,s})^2}\mathbf{E}\left[\left(\sum_{m,s} U_m(t)h_{m,s}\right)^2|\mathbf{Q}(t)\right] \qquad (23) \\
&\leq K\frac{\max_m(\sum_s r_{m,s})^2}{\max_{j,k}(\sum_s r_{j,k,s})^2}\mathbf{E}\left[\sum_m U_m^2\right] \\
&< \infty.
\end{aligned}$$

Eq. (23) comes from the task allocation constraint and the last step follows as availability processes have bounded second moment.

Consider the last term of (22), as C_2 plus this is the upper bound for Lyapunov drift $\mathbf{E}[L(\mathbf{Q}(t+1))|\mathbf{Q}(t)] - L(\mathbf{Q}(t))$. Then by Lem. 2 and the fact that $\{A_j(t)\}$ are i.i.d.,

$$\begin{aligned}
&\mathbf{E}\left[\sum_j\left(l_{j,1}Q_{j,1}(A_{j,1} - S^*_{j,1}) + \sum_k l_{j,k}\left(Q_{j,k}(t)S^*_{j,p_j(k)}(t) - Q_{j,k}(t)S^*_{j,k}(t)\right)\right)|\mathbf{Q}(t)\right] \\
&= \sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))\lambda_j - \mathbf{E}\left[\sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))S^*_{j,k}|\mathbf{Q}(t)\right].
\end{aligned}$$

Note that for any $\mathbf{Q}(t)$ and $\mathbf{U}(t)$:

$$\sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))S^*_{j,k} \geq \max_{\mathbf{a}\in C(\mathbf{U}(t))}\sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))a_{j,k}.$$

Note that in the optimal allocation $\{S^*_{j,k}\}$, $S^*_{j,k} \geq 0$ only if $\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))) \geq 0$ (otherwise, just setting them to 0 gives a better allocation). So,

$$\sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))S^*_{j,k} \geq \max_{\mathbf{a}\in C(\mathbf{U}(t))}\sum_j\sum_{k=1}^{K_j}|\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k} - Q_{j,r})|^+a_{j,k}.$$

Hence,

$$\begin{aligned}
&\mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} \sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t)-Q_{j,r}(t))S^*_{j,k}|\mathbf{Q}(t)\right]\\
&=\mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}Q_{j,k}-Q_{j,r}|^+ S^*_{j,k}|\mathbf{Q}(t)\right]\\
&\geq \sup_{\mathbf{a}\in\mathbf{C}} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}-Q_{j,r})|^+ a_{j,k}|\mathbf{Q}(t)\right]\\
&\geq \sup_{\substack{\mathbf{a}\in\mathbf{C}:\\ a_{j,k}=a_j\, 1\leq k\leq K_j}} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}Q_{j,k}-Q_{j,r}|^+ a_{j,k}|\mathbf{Q}(t)\right]\\
&=\sup_{\mathbf{a}\in\mathcal{C}} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}Q_{j,k}-Q_{j,r}|^+ a_j|\mathbf{Q}(t)\right]\\
&\geq \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}Q_{j,k}-Q_{j,r}|^+ \lambda_j|\mathbf{Q}(t)\right]+\epsilon\sum_j\sum_{k=1}^{K_j}|\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}-Q_{j,r})|^+,
\end{aligned}$$

because $\boldsymbol{\lambda}+\epsilon\in\mathcal{C}$.

As,

$$\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}Q_{j,k}-Q_{j,r}|^+\lambda_j \geq \sum_j\sum_{k=1}^{K_j}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t)-Q_{j,r}(t))\lambda_j$$

we have

$$\mathbf{E}[L(\mathbf{Q}(t+1))-L(\mathbf{Q}(t))|\mathbf{Q}(t)]\leq C_2-\epsilon\sum_j\sum_{k=1}^{K_j}|\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}-Q_{j,r})|^+.$$

Note that for $\{Q_{j,k}\}$ sufficiently large, $\sum_j\sum_{k=1}^{K_j}|\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t)-Q_{j,r}(t))|^+$ is also large. This is because if $\{Q_{j,k}\}$ is larger than B (in max-norm) then there exists a j such that $\max_k Q_{j,k}>B$. Now consider two cases, if $Q_{j,L_j}\geq\frac{B}{2}$ for some leaf node L_j then we have the drift $\leq C_2-\epsilon\frac{B}{2}$ which can be made strictly negative by choosing B appropriately.

If $Q_{j,L_j}<\frac{B}{2}$ for all leaf nodes, then there exists a k_0 such that $Q_{j,k_0}>B$. Note the following for the set of nodes T_{k_0} in the subtree rooted at k_0 and L_{k_0} being the leaves of T_{k_0}:

$$\sum_{k\in T_{k_0}}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}-Q_{j,r})=\sum_{l\in L_{k_0}}(Q_{j,k_0}-Q_{j,l}). \tag{24}$$

Hence, we have that $\sum_{k\in T_{k_0}}\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}-Q_{j,r})\geq l_{j,k_0}\frac{B}{2}\geq\frac{B}{2}$.

Thus we show strictly negative drift for sufficiently large $\{Q_{j,k}\}$ and the drift is bounded by $C_2<\infty$. Hence, by the Foster-Lyapunov theorem in [18] we have that for any initial distribution, there exists a $d\in\{1,2,\ldots\}$ such that $\mathbf{Q}(td)\to\pi$ in distribution.

To prove finite expectation we consider the following.

$$\mathbf{E}[L(\mathbf{Q}(t+1)) - L(\mathbf{Q}(t))|\mathbf{Q}(t)] \leq C_2 - \epsilon \sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+$$

which implies that

$$\mathbf{E}[L(\mathbf{Q}(t+1)) - L(\mathbf{Q}(t))] \leq C_2 - \epsilon \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+\right].$$

Summing both sides from 0 to T, we get:

$$\frac{1}{T}\sum_{t=1}^{T} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r} Q_{j,k}(t) - Q_{j,r}(t)|^+\right] \leq \frac{1}{\epsilon}\left(C_2 - \frac{1}{T}\mathbf{E}[L(\mathbf{Q}(T+1))] + \mathbf{E}[L(\mathbf{Q}(0))]\right).$$

As $\mathbf{E}[L(\mathbf{Q}(0))]$ finite, for any initial condition we have

$$\frac{1}{T}\sum_{t=1}^{T} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+\right] < C_3,$$

for all T.

As all terms are positive, for any $d \in \{1, 2, \dots\}$,

$$\lim_{T\to\infty} \frac{d}{T}\sum_{t=1}^{T} \mathbf{E}\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(td) - Q_{j,r}(td))|^+\right] < C_3.$$

By the ergodicity of a Markov chain in a positive recurrent class this implies that

$$\mathbf{E}_\pi\left[\sum_j \sum_{k=1}^{K_j} |\sum_{r \in c_j(k)} l_{j,r}(Q_{j,k}(td) - Q_{j,r}(td))|^+\right] < C_3.$$

This proves that $\mathbf{E}_\pi\left[Q_{j,L_j}\right] < C_3$ for any leaf node l_j.

By (24) we have that for any $k \in T_j$,

$$\begin{aligned} l_{j,k}Q_{j,k} &= \sum_{l_j} Q_{j,l_j} + \sum_{k' \in T_k} \sum_{r \in c(k')} l_{j,r}(Q_{j,k'} - Q_{j,r}) \\ &\leq \sum_{l_j} Q_{j,l_j} + \sum_{k' \in T_k} |\sum_{r \in c(k')} l_{j,r}(Q_{j,k'} - Q_{j,r})|^+. \end{aligned}$$

Hence, it follows that $\mathbf{E}_\pi[Q_{j,k}] < \infty$. This implies that $\mathbf{E}_\pi\left[\sum_j \sum_{k=1}^{K_j} Q_{j,k}\right] < \infty$ and so the proof is complete.

C. Proof of Theorem 3

In deriving (22) we did not use any property of the allocation $\{S^*_{j,k}\}$ other than the fact that it has to satisfy the step allocation constraint. Hence, this upper bound for Lyapunov drift is valid for any arbitrary feasible allocation $\{S_{j,k}\}$.

Hence, under the LP-relaxation base allocation $\{S^R_{j,k}\}$ by Lem. 2 we have:

$$\mathbf{E}[L(\mathbf{Q}(t+1)) - L(\mathbf{Q}(t))|\mathbf{Q}(t)] \leq C_2 + 2\mathbf{E}\left[\sum_j \sum_{r \in c_j(k)} l_{j,r}\left(Q_{j,k} - Q_{j,r}\right) S^R_{j,k} |\mathbf{Q}(t)\right]. \tag{25}$$

Note that for the optimum of the problem in (7), $\{S^R_{j,k}(t)\}$, the following is true.

$$\begin{aligned}
&\sum_j \sum_{k=1}^{K_j} \sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))S^*_{j,k} \\
&= \sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+ S^*_{j,k} \\
&\leq \sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+ \hat{S}_{j,k}
\end{aligned}$$

This is because (7) solves a relaxed problem and the optimal allocation has $S^*_{j,k} = \hat{S}_{j,k} = 0$ for $\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))$. As $S^R_{j,k} = \lfloor \hat{S}_{j,k} \rfloor$ we have that

$$\sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+ S^*_{j,k} \leq \sum_j \sum_{k=1}^{K_j} |\sum_{r\in c_j(k)} l_{j,r}(Q_{j,k}(t) - Q_{j,r}(t))|^+ (S^R_{j,k} + 1)$$

Hence, for any $\boldsymbol{\lambda}$ such that $\boldsymbol{\lambda} + \mathbf{1}(1+\epsilon) \in \mathcal{C}$ using the same proof as above we can show that the system is stable.

D. Proof of Theorem 4

This proof has the following structure. As the total number of incomplete jobs is equal to the total number of unallocated steps, we first show that the total number of unallocated steps at depth 0 (i.e., at the root of each T_j) across all types have the desirable property. Then we show that this property propagates.

*Proof for Depth-*0 *Steps:*

This part is same as the proof [24] of performance guarantee of GreedyJob algorithm in [15]. We present it here for the sake of completeness.

Consider the different types of unallocated steps at depth 0. These are given by $\{Q_{j,1}(t) : j \in [N]\}$.

Consider the following processes: for each $s \in [S]$, $Q^s_1(t) = \sum_{j:r_{j,1,s}>0} Q_{j,1} r_{j,1,s}$ which represent the number of unserved hours of skills s for all steps at depth 0.

We now construct another process $\tilde{Q}_1$ such that it dominates the process $\sum_s Q^s_1$. So, if we can upper bound $\tilde{Q}_1$, then the same bound applies for $\sum_s Q^s_1$. Hence, in turn we get a bound for $\{Q_{j,1}(t)\}$ (since $\min\{r_{j,k,s} > 0\} = \Theta(1)$ by the assumption that $\{r_{j,k,s}\}$ do not scale with the system size).

To construct a suitable $\tilde{Q}_1$, we make the following observation about the dynamics of Q^s_1 and $\{Q_{j,1}\}$. At each time t, $\sum_j A_{j,1}(t) r_{j,1,s}$ amount of s skill-hour is brought to add to Q^s_1. Also, this queue gets some service depending on the available agent hours.

At time t, $\sum_m U_m(t) h_{m,s}$ s-skill hour of service is brought by the agents.

For a step to be allocated, all of its tasks must find an allocation. Hence, for a step in type j-job to find an allocation it must get $r_{j,1,s}$ hours of service from each skill s. Thus at any time t any skill s queue gets a service of at least

$$\min_{s\in[S]} \sum_m U_m h_{m,s} - \bar{r},$$

where $\bar{r} = \max\{r_{j,k,s}\}$, due to the following. For each skill, $\sum_s U_m h_{m,s}$ hours are available. Note that a step can be allocated if all its tasks find allocations, the converse of which is also true. That is, if all tasks of a step find allocation, then the step can be allocated. As $\min_{s\in[S]} \sum_s U_m h_{m,s}$ hours of service are brought by the agents for each skill, at least $\min_{s\in[S]} \sum_s U_m h_{m,s} - \bar{r}$ of s-skill hours are served (because a maximum of $\bar{r}$ can be wasted, as no task is of size more than $\bar{r}$). Note that as depth d steps have priority in Priority Greedy algorithm over steps at depth $\geq d+1$, they do not have to share resource with higher-depth steps. So at depth d, $\min_{s\in[S]} \sum_s U_m h_{m,s}$ is available for service to steps at depth $\leq d$.

Also, note that the amount of required service brought to the queue Q_1^s at time t is upper-bounded by

$$\max_{s\in[S]} \sum_j A_{j,1}(t) r_{j,1,s}.$$

Consider a process $\tilde{Q}_1^s$ with evolution:

$$\tilde{Q}_1^s(t+1) = \max(\tilde{Q}_1^s(t) + \max_{s\in[S]} \sum_j A_{j,1}(t) r_{j,1,s} - \min_{s\in[S]} \sum_m U_m h_{m,s} + \bar{r}, 0).$$

Note that given $\tilde{Q}_1^s(t_0) \geq Q_1^s(t_0)$ at some t_0, the same holds true for all $t \geq t_0$. This is because for $x, a, b \geq 0$ and $x', a', b' \geq 0$, with $x \geq x'$, $a \geq a'$ and $b \leq b'$

$$\max(x + a - b, 0) \geq \max(x' + a' - b', 0),$$

and so the monotonicity propagates over time.

To bound $\sum_s Q_1^s$, it is sufficient to bound $\sum_s \tilde{Q}_1^s(t)$. Note that each $\tilde{Q}_1^s$ has exactly the same evolution, so let us consider

$$\tilde{Q}_1 := S\tilde{Q}_1^1,$$

which bounds $\sum_s Q_1^s$.

From the evolution,

$$\tilde{Q}_1(t+1) = \max(\tilde{Q}_1(t) + S\max_{s\in[S]} \sum_j A_{j,1}(t) r_{j,1,s} - S\min_{s\in[S]} \sum_m U_m h_{m,s} + \bar{r}, 0),$$

and we can write the Loynes' construction for this process which has the same distribution as the following process (and for simplicity we use the same notation, as we are interested in the distribution):

$$\tilde{Q}_1^1(0) = \max_{\tau\leq 0} \sum_{\tau\leq t\leq 0} (S\max_{s\in[S]} \sum_j A_{j,1}(t) r_{j,1,s} - S\min_{s\in[S]} \sum_m U_m h_{m,s} + \bar{r}),$$

assuming that the process started at $-\infty$.

Let us define $X_s(t)$ and $Y_s(t)$ as follows: $X_s(t) := \sum_j A_{j,1}(t) r_{j,1,s}$ and $Y_s(t) := \sum_m U_m h_{m,s}$. Then,

$$\tilde{Q}_1^1(0) = \max_{\tau\leq 0} \sum_{\tau\leq t\leq 0} S(\max_s X_s(t) - \min_s Y_s(t) + \bar{r}).$$

Now, for any $\theta > 0$:

$$\begin{aligned}\mathbf{P}(\sum_j Q_{j,1} > \bar{r}q) &\leq \mathbf{P}(\sum_s Q_1^s > q) \\ &\leq \mathbf{P}(\tilde{Q}_1(0) > q) \\ &= \mathbf{P}(\theta\tilde{Q}_1(0) > \theta q) \\ &= \mathbf{P}(\exp(\theta\tilde{Q}_1(0)) > \exp(\theta q)) \\ &\leq \mathbf{E}[\exp(-\theta q)]\mathbf{E}[\exp(\theta\tilde{Q}_1(0))].\end{aligned} \tag{26}$$

Now,

$$\begin{aligned}\mathbf{E}[\exp(\theta\tilde{Q}_1(0))] &= \mathbf{E}\left[\exp(\theta S\left(\max_{\tau\leq 0}\sum_{\tau\leq t\leq 0}(\max_s X_s(t) - \min_s Y_s(t) + \bar{r})\right))\right] \\ &\leq \sum_{\tau\leq 0}\mathbf{E}\left[\exp(\theta S\sum_{\tau\leq t\leq 0}(\max_s X_s(t) - \min_s Y_s(t) + \bar{r}))\right],\end{aligned} \tag{27}$$

where inequality (27) follows because for any random variables $\{Z_j\}$, $\exp(\theta\mathbb{Z}_j)$ are positive random variables and the sum of positive values is more than their maximum.

Next, we bound the term within the summation over $\tau \leq 0$ in (27).

$$\mathbf{E}[\exp(\theta S\sum_{\tau\leq t\leq 0}(\max_s X_s(t) - \min_s Y_s(t) + \bar{r}))] \leq \prod_{\tau\leq t\leq 0}\mathbf{E}[\exp(\theta(\max_s X_s(t) - \min_s Y_s(t) + \bar{r})))], \tag{28}$$

which follows because $X_s(t)$, $Y_s(t)$ are i.i.d. over time.

Next we bound the term within the product in (28),

$$\mathbf{E}\left[e^{\theta S(\max_s X_s(t) - \min_s Y_s(t) + \bar{r})}\right] \leq \sum_{s,s'}\mathbf{E}\left[e^{\theta S(X_s(t) - Y_{s'}(t) + \bar{r})}\right], \tag{29}$$

where this follows for the same reason as (27).

The following lemma regarding an outer bound to the capacity region will be useful later.

Lemma 3. *Let* $\mathcal{C}_{I,F}^O = \{\boldsymbol{\lambda} : \sum_{j,k}\lambda_j r_{j,k,s} \leq \sum_m u_m h_{m,s}\}$. *Then* $\mathcal{C}_{I,F} \subset \mathcal{C}_{I,F}^O$

Proof. Consider any $\boldsymbol{\lambda} \in \mathcal{C}_{I,F}$. Then by definition of $\mathcal{C}_{I,F}$, $\boldsymbol{\lambda}^E \in \mathbf{C}$ and there exists $\mathbf{c}(\mathbf{u}) \in convC_{I,F}(\mathbf{u})$ such that

$$\boldsymbol{\lambda}^E \leq \sum_{\mathbf{u}}\Gamma(\mathbf{u})\mathbf{c}(\mathbf{u}).$$

Also, note that for each $\mathbf{c}(\mathbf{u}) \in convC_{I,F}(\mathbf{u})$, there exists an $\{\alpha_v(\mathbf{u}) \geq 0, 1 \leq v \leq V_{\mathbf{u}} : \sum_{v=1}^{V_{\mathbf{u}}}\alpha_k = 1\}$ and $a_v(\mathbf{u}) \in C_{I,F}(\mathbf{u}), 1 \leq v \leq V_{\mathbf{u}}$ such that

$$\sum_{v=1}^{V_{\mathbf{u}}} a_k(\mathbf{u})\alpha_v(\mathbf{u}) = \mathbf{c}(\mathbf{u}).$$

Note that as $a_v(\mathbf{u}) \in C_{I,F}(\mathbf{u})$ so by allocation constraint

$$\sum_{j,k} a_{v,j,k}(\mathbf{u})r_{j,k,s} \leq \sum_m u_m h_{m,s},$$

where $a_{v,j,k}$ are the number of (j,k) steps that have been allocated under $a_v(\mathbf{u})$ allocation. This in turn implies that:

$$\begin{aligned}
&\sum_{j,k} a_{v,j,k}(\mathbf{u})\alpha_v(\mathbf{u})r_{j,k,s} \leq \sum_m u_m h_{m,s} \\
&\implies \sum_{j,k} \mathbf{c}_{j,k} r_{j,k,s} \leq \sum_m u_m h_{m,s} \\
&\implies \sum_{j,k} \lambda_j r_{j,k,s} \leq \sum_m u_m h_{m,s}.
\end{aligned}$$

Hence, $\boldsymbol{\lambda} \in \mathcal{C}_{I,F}$ implies that $\boldsymbol{\lambda} \in \mathcal{C}_{I,F}^O$. □

Let $\mathcal{A}_j(\theta) = \mathbf{E}\left[e^{\theta A_j(t)}\right]$ and $\mathcal{U}_m(\theta) = \mathbf{E}\left[e^{\theta U_m(t)}\right]$ for $j \in [N]$ and $m \in [M]$. For $\theta \in \mathbb{R}$, then,

$$\begin{aligned}
\mathbf{E}\left[e^{\theta(X_s(t)-Y_{s'}(t))}\right] &= \mathbf{E}\left[e^{\theta X_s(t)}\right]\mathbf{E}\left[e^{-\theta Y_{s'}(t)}\right] \\
&= \mathbf{E}\left[e^{\theta\sum_j A_j(t)r_{j,1,s}+\bar{r}}\right]\mathbf{E}\left[e^{-\theta\sum_i U_i(t)h_{i,s}}\right] \\
&= e^{\theta\bar{r}}\prod_j \mathbf{E}\left[e^{\theta A_j(t)r_{j,1,s}}\right]\prod_i \mathbf{E}\left[e^{-\theta U_i(t)h_{i,s'}}\right] \\
&= e^{\theta\bar{r}}\prod_j \mathcal{A}_j(\theta r_{j,1,s})\prod_i \mathcal{U}_i(-\theta h_{i,s'}) \\
&= \exp\left(\theta\bar{r} + \sum_j \log\mathcal{A}_j(\theta r_{j,1,s}) + \sum_i \log\mathcal{U}_i(-\theta h_{i,s'})\right).
\end{aligned} \tag{30}$$

Note that as $\boldsymbol{\lambda} \in \alpha\mathcal{C}$, by the definition of $\mathcal{C}_{I,F}^O$, $\sum_j \lambda_j r_{j,1,s} < \alpha\sum_m \mu_m h_{m,s}$ and by assumption $|\sum_m \mu_m h_{m,s} - \sum_m \mu_m h_{m,s'}| \leq \text{subpoly}(N)$ which is used in the following.

First consider the Gaussian-dominated case. Since the process variance is no more than the mean and the moment generating function of the variance is upper-bounded by that of a zero-mean Gaussian:

$$\begin{aligned}
\log\mathcal{A}_j(\theta r_{j,s}) &\leq \lambda_j\theta r_{j,1,s} + \lambda_j\frac{(\theta r_{j,1,s})^2}{2} \\
\log\mathcal{U}_i(-\theta h_{i,s}) &\leq -\mu_j\theta h_{i,s} + \mu_j\frac{(\theta h_{i,s})^2}{2}.
\end{aligned}$$

Note that for any two functions k_1x^2 and k_2x, $\lim_{x\to 0} k_2x/k_1x^2 = \infty$, and hence for any $\epsilon \in (0,1)$ there exists $x^* > 0$ such that for all $x < x^*$, $k_1x^2/k_2x < \epsilon$. Hence for any $\epsilon \in (0,1)$, there exist $\theta_{j,s}^*, \theta_{i,s}^* > 0$, for all i,j,s such that for all $\theta < \theta^* = \min_{i,j,s}(\theta_{j,s}^*, \theta_{i,s}^*)$,

$$\log\mathcal{A}_j(\theta r_{j,1,s}) \leq \lambda_j\theta^* r_{j,1,s}(1+\epsilon) \tag{31}$$

$$\log\mathcal{U}_i(-\theta h_{i,s}) \leq -\mu_i\theta h_{i,s}(1-\epsilon). \tag{32}$$

Note that since N, S, and M are finite and $\theta_{j,s}^*, \theta_{i,s}^* > 0$, for all $i,j,s,\theta^* > 0$. Moreover, note that θ^* does not depend on $\boldsymbol{\lambda}, \boldsymbol{\mu}$ since the ratio of the linear and quadratic terms in the log moment generating functions are independent of $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$.

As $e^\theta - 1 = \sum_{k=1}^{\infty} \frac{\theta^k}{k!}$, for the Poisson-dominated case we have:

$$\log \mathcal{A}_j(\theta r_{j,s}) \leq \lambda_j \sum_k \frac{(\theta r_{j,1,s})^k}{k!}$$

$$\log \mathcal{U}_i(-\theta h_{i,s}) \leq \mu_j \sum_k \frac{(-\theta h_{i,s})^k}{k!}.$$

Again, by the same argument, we can have a θ^* for which (31) and (32) are satisfied. As $|\sum_i \mu_i h_{i,s} - \sum_i \mu_i h_{i,s'}| = o(N^\delta), \forall \delta > 0$, and $\sum_i \mu_i h_{i,s} = \Omega(N^c), c > 0$, for all $\theta < \theta^*$ we have:

$$\begin{aligned}\mathbf{E}\left[e^{\theta(X_s(t)-Y_{s'}(t))}\right] &\leq \exp\left(\theta^*\bar{r} + \sum_j \lambda_j \theta^* r_{j,1,s}(1+\epsilon) - \sum_i \mu_i \theta h_{i,s}(1-\epsilon) + \theta^* o(\mu_i \theta h_{i,s})\right) \\ &\leq \exp\left(\theta\left(\bar{r} - \sum_i \mu_i h_{i,s}(\alpha - 2\epsilon)\right)\right). \end{aligned} \tag{33}$$

Note (33) follows from the fact $\lambda \in (1-\alpha)\mathcal{C}^O$. As $\epsilon > 0$ can be chosen arbitrarily small, we can have $\alpha - 2\epsilon > 0$. Since $\sum_i \mu_i h_{i,s} > \sum_j (1-\alpha)\lambda_j r_{j,s}$ and $\sum_j \lambda_j r_{j,1,s}$ scales with $\lambda(N)$, for sufficiently large λ_α with $\lambda_j \geq \lambda_\alpha$ for all j, we have $\bar{r} - \sum_i \mu_i h_{i,s}(\alpha - \epsilon) \leq -\gamma \sum_i \mu_i h_{i,s}(\alpha - \epsilon)$, for some $\gamma > 0$. Thus, we have for some $\theta > 0$,

$$\mathbf{E}\left[e^{\theta S(X_s(t)-Y_{s'}(t))}\right] \leq \exp\left(-\theta S K(N)\right), \tag{34}$$

where $K(N)$ scales with N no slower than $\sum_{s:r_{j,1,s}>0} \lambda_j(N) = \Omega(N^c)$, $c > 0$.

Thus,

$$\mathbf{E}\left[e^{\theta S(\max_s X_s(t) - \min_s Y_s(t))}\right] \leq S^2 \exp\left(-\theta S K(N)\right).$$

Hence, from (27), (28), and (29) we have that

$$\begin{aligned}\mathbf{E}[\exp(\theta^* \sum_s \tilde{Q}_1^s(0))] &= \mathbf{E}[\exp(\theta^* S \tilde{Q}_1^1(0))] \\ &= \mathbf{E}[\exp(\theta^* S \tilde{Q}_1(0))] \\ &\leq \sum_{\tau \leq 0} S^{2|\tau|} \exp(-\theta^* S K(N)|\tau|) \\ &\leq c',\end{aligned}$$

because $S^2 < \exp(\theta^* S K(N))$ for sufficiently large N.

Note that though we proved $\mathbf{E}[\exp(\theta Q_1(t))] < c'$ for $t = 0$, this holds for any finite t (exactly the same proof). This in turn implies that the number of unallocated steps in depth 0 have bounded exponential moment for some $\theta > 0$. This will be used in the remainder of the proof where we show that the same is true for all depths.

Induction over Depths, d to $d+1$:

Now we show that if the total number of unallocated steps at depth d satisfies $\mathbf{E}[\exp(\theta Q(0))] < c'$, then the same is true for $d+1$. To show the same result for steps at all depths we consider the following process. Let $d_j(k)$ be the depth of k in T_j, then $Q_{d+1}^s(t) = \sum_{j,k:d_j(k)\leq d+1} Q_{j,k} r_{j,k,s}$ represents the number of unserved hours of skills s for all steps in the system.

Like in the case of the proof for depth 0, we construct process $\tilde{Q}^s_{d+1}$ such that $\sum_s \tilde{Q}^s_{d+1}$ dominates the process $\sum_s Q^s_{d+1}$. Using the same argument as before, at any time t any skill s queue gets a service of at least

$$\min_{s\in[S]} \sum_m U_m h_{m,s} - \bar{r},$$

and the amount of required service brought to the queue Q^s at time t is upper-bounded by

$$\max_{s\in[S]} \sum_{j,k:d_j(k)\le d+1} A_{j,k}(t) r_{j,k,s}.$$

Then using the same argument, the process

$$\tilde{Q}_{d+1}(t+1) = \left| (\tilde{Q}_{d+1}(t) + S \max_{s\in[S]} \sum_{j,k:d_j(k)\le d+1} A_{j,k}(t) r_{j,k,s} - \min_{s\in[S]} \sum_m U_m h_{m,s} + \bar{r} \right|^+$$

upper-bounds the process $\sum_s Q^s_{d+1}$. Then we can follow the steps that we followed using X_s and Y_s previously. Let $X'_s := \sum_{j,k:d_j(k)\le d+1} A_{j,k}(t) r_{j,k,s}$ and $Y'_s := \sum_m U_m h_{m,s}$, respectively. But note that $A_{j,k}$ for $k > 1$ is not an external i.i.d. process, rather it is the number of steps of type $A_{j,p_j(k)}$ that were completed. Hence, we cannot follows the exactly same steps. Note that

$$\begin{aligned}
\mathbf{E}[\exp(\theta S \sum_{\tau\le t\le 0} (\max_s X'_s(t) - \min_s Y'_s(t) + \bar{r}))] &\le \mathbf{E}[\exp(\theta \sum_{\tau\le t\le 0} \max_{s,s'\in[S]} (X'_s(t) - Y'_{s'}(t) + \bar{r}))] \\
&\le \mathbf{E}[\sum_{s,s'\in[S]} \exp(\theta \sum_{\tau\le t\le 0} (X'_s(t) - Y'_{s'}(t) + \bar{r}))] \\
&= \sum_{s,s'\in[S]} \mathbf{E}[\exp(\theta S \sum_{\tau\le t\le 0} (X'_s(t) - Y'_{s'}(t) + \bar{r}))].
\end{aligned}$$

Also note that,

$$\sum_{\tau\le 0} \mathbf{E}[\exp(\theta S \sum_{\tau\le t\le 0} \max_{s,s'\in[S]} (X'_s(t) - Y'_{s'}(t) + \bar{r}))] \le \sum_{s,s'\in[S]} \sum_{\tau\le t\le 0} \mathbf{E}[\exp(\theta S \sum_{\tau\le t\le 0} (X'_s(t) - Y'_{s'}(t) + \bar{r}))].$$

So, we investigate $\mathbf{E}[\exp(\theta \sum_{\tau\le t\le 0}(X'_s(t) - Y'_{s'}(t) + \bar{r}))]$.

$$\begin{aligned}
&\mathbf{E}[\exp(\theta \sum_{\tau\le t\le 0} (X'_s(t) - Y'_{s'}(t) + \bar{r}))] \\
&= \exp(\bar{r}\theta)\mathbf{E}[\exp(\theta \sum_{\tau\le t\le 0} (\sum_{j,k} A_{j,k}(t) r_{j,k,s} - \sum_m U_m(t) h_{m,s'}))] \\
&= \exp(\bar{r}\theta)\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\le d+1} r_{j,k,s} \sum_{\tau\le t\le 0} A_{j,k}(t) - \sum_m \sum_{\tau\le t\le 0} U_m(t) h_{m,s'}))].
\end{aligned}$$

Note that $\sum_{\tau\le t\le 0} A_{j,k}(t)$ represent the creation (or appearance/arrival) of steps of type (j,k) between time τ and 0 (with a similar interpretation for agents in case of $\sum_{\tau\le t\le 0} U_m(t)$), which we denote by $A_{j,k}(\tau:0)$ (and $U_m(\tau:0)$), respectively.

Now there is an important observation about $A_{j,k}(\tau:0)$:

$$A_{j,k}(\tau:0) \le Q_{j,p_j(k)}(\tau-1) + A_{j,p_j(k)}(\tau-1:-1), \tag{35}$$

where $p_j(k)$ is the parent of k in T_j, due to the following. As each job takes one slot to be served, no job whose step $(j, p_j(k))$ completed after -1 can have its step (j, k) be available for service at or before 0. Thus by induction on the function p_j we can write

$$A_{j,k}(\tau:0) \leq \sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) + A_{j,1}(\tau - d_j(k) : -d_j(k)), \tag{36}$$

as $d_j(k) = d + 1$ by the inductive assumption. Note that from 1 to k (at depth $d + 1$) there is a unique d-length path and hence, on that path w.l.o.g. we denote the respective steps by (j, w) where w is its depth on that path.

Hence,

$$\begin{aligned}
&\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s} \sum_{\tau\leq t\leq 0} A_{j,k}(t) - \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'}))] \\
&\leq \mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}(\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) \\
&\quad + A_{j,1}(\tau - d_j(k) : -d_j(k))) - \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'}))]
\end{aligned}$$

Note that $\sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'}$ is independent of

$$\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}\left(\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) + A_{j,1}(\tau - d_j(k) : -d_j(k))\right),$$

because $A_{j,1}$ are i.i.d. (independent of U_m) and $Q_{j,w}(\tau - d + w)$ does not depend on $U_m(\tau : 0)$ for $d \geq w \geq 1$. Hence,

$$\begin{aligned}
&\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s} \sum_{\tau\leq t\leq 0} A_{j,k}(t) - \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'}))] \\
&\leq \mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}\left(\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) + A_{j,1}(\tau - d_j(k) : -d_j(k))\right)))] \\
&\quad \times \mathbf{E}[\exp(-\theta \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'})].
\end{aligned}$$

We use the previously derived bound for $\mathbf{E}[\exp(-\theta \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'})]$. So, we only concern ourselves with

$$\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}\left(\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) + A_{j,1}(\tau - d_j(k) : -d_j(k))\right)))].$$

Consider any $Q_{j,w}(\tau - 1 - d + w)$ at depth w, then $A_{j,1}(\tau - d - 1)$ is independent of it. As $A_{j,1}$ are i.i.d. and future arrivals in a queue are independent of present and past queue-lengths, we have

$$\begin{aligned}
&\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}(\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w) + A_{j,1}(\tau - d_j(k) : -d_j(k)))] \\
&= \mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s} \sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w)))] \times \mathbf{E}[\exp(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s} A_{j,1}(\tau - d_j(k) : -d_j(k)))]
\end{aligned}$$

For the second term we obtain a bound using previous techniques and note that since $\boldsymbol{\lambda} \in \alpha\mathcal{C}$,

$$\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}\mathbf{E}[A_{j,1}] \leq \sum_m \mu_m h_{m,s'},$$

which in the same way as above will imply that for some $K(N)$ and some $\theta > 0$,

$$\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}A_{j,1}(\tau - d_j(k) : -d_j(k)) - \sum_m \sum_{\tau\leq t\leq 0} U_m(t)h_{m,s'}))]$$
$$\leq \exp(-\theta K(N)\tau).$$

Note that

$$\mathbf{E}[\exp(\theta(\sum_{j,k:d_j(k)\leq d+1} r_{j,k,s}\sum_{w=1}^{d} Q_{j,w}(\tau - 1 - d + w)))] < \infty$$

by the inductive assumption that the number of unallocated steps at depth $\leq d$ have finite exponential moments.

So we have that

$$\mathbf{E}[\exp(\theta \sum_{\tau\leq t\leq 0} (X_s(t) - Y_s(t) + \bar{r}))] < c_1 \exp(-\theta K(N)\tau),$$

and so, in turn (using the same steps as above) Q^s has finite exponential moment for some θ. The rest of the steps are similar to above and we get the desired result that

$$\mathbf{E}[\exp(\theta \sum_{j,k:d_j(k)\leq d+1} Q_{j,k})] < \infty.$$

By induction on d, we have proven that the total number of unallocated steps over all types of jobs have finite exponential moment (say c').

Therefore,

$$\mathbf{P}(\sum_{j,k} Q_{j,k} > q) \leq \exp(-\theta q)\mathbf{E}[\exp(\theta \sum_{j,k} Q_{j,k})]$$
$$\leq c' \exp(-\theta q).$$

So for $q = \frac{3\log N}{\theta}$, we have the result (as c' is constant).

E. Proof of Theorem 5

The following lemma is useful for the proof.

Lemma 4. *Any feasible solution of problem* (9) *is a feasible solution of problem* (10)*.*

Proof: Among the total R available agents, let R_m be of type m, and let them be denoted $i_1, i_2, \ldots, i_{R_m}$. For an allocation $\mathbf{a}$ in the formulation (9), let $\tilde{h}_{i_k,s}$ be the time that agent i_1 is assigned for skill s. Then among all these type m agents, the total contribution to skill s is $\sum_{i=1}^{R_m} \tilde{h}_{i_k,s}$.

Note that the total time for these R_m agents is $R_m h_m$. Now by the allocation constraint we have $\sum_s \sum_{i=1}^{R_m} \tilde{h}_{i_k,s} \leq R_m h_m$. If we choose $\alpha_{m,s} \geq \sum_{i=1}^{R_m} \tilde{h}_{i_k,s}/R_m h_m$, this is a valid allocation in formulation (10), as $\sum_s \alpha_{m,s} \leq 1$ and it also meets the allocation constraint. So, for this $\alpha_{m,s}$ the allocation $\mathbf{a}$ is a valid allocation in problem (10).

This proves that every valid allocation in (9) is also a valid allocation in (10). ■

Since solving (10) yields a feasible allocation, the lemma implies the two problems are actually alternate formulations of one another. The rest of the proof follows the same steps as the proof of Thm. 3.

F. Proof of Theorem 6

The result can be derived in the same way as Thm. 4, through the use of the following lemmata.

Lemma 5. *Let $\{R_i\}$ be i.i.d. Bernoulli random variables with $\mathbf{P}(R_i = 1) = p \in (0, 1)$ and N be a random variable independent of $\{R_i\}$ with a moment generating function $M_N(\theta)$. Then*

$$\mathbf{E}[\exp(\theta \sum_{i=1}^{N} R_i)] = M_N(\log(p \exp(\theta) + (1 - p))).$$

Proof:

$$\begin{aligned}
\mathbf{E}[\exp(\theta \sum_{i=1}^{N} R_i)] &= \mathbf{E}\left[\mathbf{E}[\exp(\theta \sum_{i=1}^{N} R_i)|N]\right] \\
&= \mathbf{E}\left[(p \exp(\theta) + (1 - p))^N\right] \\
&= \mathbf{E}\left[\exp(\log(p \exp(\theta) + (1 - p))\, N)\right] \\
&= M_N(\log(p \exp(\theta) + (1 - p)))
\end{aligned}$$

■

Lemma 6. *Let $\mathcal{C}^O_{F,F} = \{\boldsymbol{\lambda} : \exists b_m \in [0, 1]^S$ for all m, s.t. $\sum_s b_{m,s} \leq 1, b_{m,s} > 0$ only if $s \in S_m, \sum_{m:s\in S_m} b_{m,s} h_m \mu_m > \sum_{j,k} \lambda_j r_{j,k,s}$ for all $s\}$. Then, $\mathcal{C}_{F,F} \subset \mathcal{C}^O_{F,F}$.*

Proof: This follows from the constraints in (10). Because constraints in (10) are per sample realization, and the above constraints are in expectation. So for constraints in (10) to be satisfied, the above constraints must be satisfied. ■

Lemma 7. *For any τ, $\mathbf{P}\{\cap_{t=\tau}^{\infty} \cap_s \{\psi_{m,s}(t) = p_{m,s}\}\} = 1$ such that $p_{m,s}$ solves* (12)*, assuming ties between multiple solutions are broken deterministically.*[2]

Proof: For the choice of $\gamma(t)$ (and $t_0 = -\infty$) it follows from the convergence of stochastic approximation update equations [25] and the facts that $\boldsymbol{\lambda} \in \alpha\mathcal{C}$ and $t_0 = -\infty$.

By stochastic approximation updates $\bar{A}$ and $\bar{U}$ converges almost surely to $\boldsymbol{\lambda}$ and μ respectively at any finite τ. The rest follows from Lem. 6. ■

Let $\mathcal{B}_{s,m}$ be the set of agents of type m that has been put into $\mathcal{B}_s$. $\mathcal{B}_{s,m}$ is Bernoulli sampling from U_m agents with probability $\psi_{m,s}(t)$.

[2]Extends to random tie breaking also, but involves more details.

We follow the same steps as in the proof of Thm. 4. Consider the work and service time brought at time t (as before). Note that work brought for skill s is $\sum_j A_{j,1}(t) r_{j,1,s}$ and service time brought by agents for skill s is $\sum_{m:s\in S_m} h_m \mathcal{B}_{s,m}(t)$. Hence, we can construct a queue $\tilde{Q}_1$ (as before)

$$Q_1^s(t+1) = (Q_1^s(t) + S\max_s \sum_j A_{j,1}(t) r_{j,1,s} - S\min_s \sum_{m:s\in S_m} h_m \mathcal{B}_{s,m}(t))^+$$

Following the same steps to obtain (27) we can have

$$\mathbf{E}[\exp(\theta \tilde{Q}_1)] \leq \sum_{\tau\leq 0} \mathbf{E}[\exp(\theta S \sum_{\tau\leq t\leq 0} (\max_s \sum_j A_{j,1}(t) r_{j,1,s} - \min_s \sum_{m:s\in S_m} h_m \mathcal{B}_{s,m}(t))]$$

But a result like (27) does not follow immediately, since $\mathcal{B}_{s,m}(t)$ are not independent over time and $\mathcal{B}_{s,m}(t)$ depends on $A_j(t)$ via $\{\psi_s(t)\}$.

Consider the following. Let $X_s(t) = \sum_j A_{j,1}(t) r_{j,1,s}$ and $Y_s(t) = \sum_{m:s\in S_m} h_m \mathcal{B}_{s,m}(t)$, then

$$\begin{aligned}
&\mathbf{E}[\exp(\theta S \sum_{\tau\leq t\leq 0} (\max_s X_s(t) - \min_s Y_s(t))] \\
&\quad \leq \mathbf{E}[\mathbf{E}[\exp(\theta S \sum_{\tau\leq t\leq 0} (\max_s X_s(t) - \min_s Y_s(t))|\{\psi_{m,s}(t'), s, mt' \geq \tau\}]]
\end{aligned}$$

Now by Lem. 7, for any finite τ_1,

$$\mathbf{P}\{\{\psi_s(t') = p_s, \forall t' \geq \tau_1, \forall s, m\} = 1.$$

Hence, for any finite τ_1, τ_2, $\{\psi_{m,s}(t'), s, mt' \geq \tau_1\}$ are independent of $\{A_{j,1}(t) : t\tau_2\}$.

Also, by the above argument, $X_s(t)$ and $Y_s(t)$ are independent of each other, given $\{\psi_{m,s}(t'), s, mt' \geq \tau\}$ and they are also independent over time.

$$\begin{aligned}
&\mathbf{E}[\prod_{\tau\leq t\leq 0} \mathbf{E}[\exp(\theta S(\max_s X_s(t) - \min_s Y_s(t))|\{\psi_{m,s}(t'), s, mt' \geq \tau\})]] \\
&= \mathbf{E}[\sum_{\tau\leq t\leq 0} \mathbf{E}[\exp(\theta S \max_{s,s'}(X_s(t) - Y_{s'}(t))|\{\psi_{m,s}(t'), s, mt' \geq \tau\})]] \\
&= \mathbf{E}[\sum_{\tau\leq t\leq 0} \sum_{s,s'} \mathbf{E}[\exp(\theta S(X_s(t) - Y_{s'}(t))|\{\psi_{m,s}(t'), s, mt' \geq \tau\})]] \\
&= \sum_{\tau\leq t\leq 0} \sum_{s,s'} \mathbf{E}[\mathbf{E}[\exp(\theta S X_s(t)|\{\psi_{m,s}(t'), s, mt' \geq \tau\})] \\
&\quad \times \mathbf{E}[\exp(-\theta S Y_{s'}(t)|\{\psi_{m,s}(t'), s, mt' \geq \tau\})]] \\
&= \sum_{\tau\leq t\leq 0} \sum_{s,s'} \mathbf{E}[\exp(\theta S X_s(t)|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})] \\
&\quad \times \mathbf{E}[\exp(-\theta S Y_{s'}(t)|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})] \\
&= \sum_{\tau\leq t\leq 0} \sum_{s,s'} \mathbf{E}[\exp(\theta S X_s(t)] \times \mathbf{E}[\exp(-\theta S Y_{s'}(t)|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})]
\end{aligned}$$

The first equality follows because $\max_{x\in\mathcal{X}} f(x) - \min_{x\in\mathcal{X}} g(x) = \max_{x,x'\in\mathcal{X}}(f(x) - g(x))$, for finite $\mathcal{X}$. The second equality follows due to independence of $X_s(t) - Y_s(t)$ from $t \geq \tau$ which is due to Lem. 7. The third equality

follows due to independence of $X_s(t)$ and $Y_s(t)$ given $\{\psi_s(t)\}$. The fourth equality is again due to Lem. 7, as $\{\psi_s(t') = p_s, s, mt' \geq \tau\}$ is an almost sure event. The last equality follows because $A_j(t), t \geq \tau_1$ are independent of $\{\psi_s(t') = p_s, s, mt' \geq \tau_2\}$ for any finite τ_1 and τ_2.

Note that $\mathbf{E}[\exp(\theta S X_s(t)|]$ can be evaluated exactly as in the proof of Thm. 4.

Consider for $\tau \leq t \leq 0$,

$$\begin{aligned}
&\mathbf{E}[\exp(-\theta S Y_{s'}(t)|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})] \\
&= \mathbf{E}[\exp(-\theta S \sum_{m:s\in S_m} h_m \mathcal{B}_{s,m}|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})] \\
&= \prod_{m:s\in S_m} \mathbf{E}[\exp(-\theta S h_m \mathcal{B}_{s,m}|\{\psi_{m,s}(t') = \psi_{m,s}(\tau), s, mt' \geq \tau\})]. \qquad (37)
\end{aligned}$$

Now $\mathcal{B}_{s,m}$ is Bernoulli sampling of U_m agents with probability $\psi_{m,s}(t') = \psi_{m,s}(\tau) := p_{m,s}$.

Then, by Lem. 5, for a $\tilde{\theta}_{m,s} = \log(\exp(\theta)p_{m,s} + 1 - p_{m,s})$:

$$\mathbf{E}[\exp(-\theta S h_m \mathcal{B}_{s,m}|\{\psi_s(t') = p_s, s, mt' \geq \tau\})] = \mathbf{E}[\exp(-\tilde{\theta}_{m,s} h_m U_m)].$$

Now following the same steps as in the proof of Thm. 4, we can show that (for Poisson and Gaussian dominated cases) for a sufficiently large N_α, for all $N \geq N_\alpha$, and $\delta < \frac{\alpha}{2}$,

$$\mathbf{E}[\exp(-\tilde{\theta}_{m,s} h_m U_m)] \leq \exp(-(1-\delta)\tilde{\theta}_{m,s} h_m \mu_m).$$

Now, by concavity of logarithms,

$$\begin{aligned}
\tilde{\theta}_{m,s} &= \log\left(p_{m,s}\exp(\theta) + (1 - p_{m,s})\right) \\
&\geq p_{m,s}\theta.
\end{aligned}$$

This implies

$$\mathbf{E}[\exp(-\tilde{\theta}_{m,s} h_m U_m)] \leq \exp(-(1-\delta)\theta\ p_{m,s} h_m \mu_m). \qquad (38)$$

As $\boldsymbol{\lambda} \in \alpha\mathcal{C}_{F,F}$, by Lem. 6 and 7, we have $\sum_{m:s\in S_m} p_{m,s} h_m \mu_m > (1-\alpha)\sum_{j,k} \lambda_j r_{j,k,s}$ for all s and $\epsilon < 1-\alpha$. Also, as by assumption $|\sum_{j,k} \lambda_j r_{j,k,s} - \sum_{j,k} \lambda_j r_{j,k,s'}|$ is sub-poly(N), hence

$$\sum_{m:s\in S_m} p_{s',m} h_m \mu_m > (1-\alpha)\sum_{j,k} \lambda_j r_{j,k,s}, \forall s, s'$$

This along with (38) and (37) gives the final result by following the same steps as the proof of Thm. 4.

G. Proof of Theorem 7

The assumption $\{S_m : m \in [M]\}$ is a partition implies that there exists a partition of $[S]$, say $\{\mathcal{K}_l : 1 \leq l \leq L\}$ such that for $m \in [M]$, $S_m = \mathcal{K}_l$ for some $l \in [L]$. Note that $L \leq S$.

As $\boldsymbol{\lambda} \in \mathcal{C}_{F,I}$, for any step (j,k) with $\lambda_j > 0$, the set of required skills is a subset of some $\mathcal{K}_l$. Otherwise, due to inflexibility of the steps, that step can never be allocated which contradicts that $\boldsymbol{\lambda} \in \mathcal{C}_{F,I}$.

Lemma 8. *Let* $\mathcal{C}^O_{F,I} = \{\boldsymbol{\lambda} :$ *for all* l $\sum_{j:(j,k)-$*skills* $\subset \mathcal{K}_l} \lambda_j < \sum_{m:S_m=\mathcal{K}_l} \mu_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor\}$. *Then, under the conditions in Thm. 7,* $\mathcal{C}_{F,I} \subset \mathcal{C}^O_{F,I}$.

Proof: Follows by noting the fact that under the conditions in Thm. 7 the steps with skill requirements in $\mathcal{K}_l$ has can only be served by agents of type m with $S_m = \mathcal{K}_l$. Also, note that agents with skills in $\mathcal{K}_l$ cannot serve any other kinds of steps.

Also, as each step is of same size $r_{1,1,1}$ a type m agent can serve at most $\lfloor \frac{h_m}{r_{1,1,1}} \rfloor$ steps with skill requirements in S_m.

This implies that if the agent availability is $\mathbf{u}$ then $a_{j,k}$ steps can be served only if

$$\sum_{j:(j,k)-\text{skills } \subset \mathcal{K}_l} a_{j,k} < \sum_{m:S_m=\mathcal{K}_l} u_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor.$$

The rest follows because the sample path constraint is true only if the constraint is true in expectation. ■

Notice that the condition

$$\sum_{j:(j,k)-\text{skills } \subset \mathcal{K}_l} \lambda_j < \sum_{m:S_m=\mathcal{K}_l} \mu_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor\}$$

can be written as

$$\sum_d \sum_{j:(j,k)\text{-skill } \subset \mathcal{K}_l \cap \mathcal{A}_d} \lambda_j < \sum_{m:S_m=\mathcal{K}_l} \mu_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor.$$

This will be useful later.

Coming back to the main proof, we consider a queue for each $l \in [L]$ and depth d ($\leq D$), $\tilde{Q}^l_d(t)$. This queue represents the number of unallocated steps with skill requirements in $\mathcal{K}_l$ that are at depth $\leq d$ on respective precedence trees.

Note that such steps with skill requirements in $\mathcal{K}_l$ can only be served by agents of types $\{m : S_m = \mathcal{K}_l\}$. Note that just as in Restricted Greedy, steps at depth $\leq d$ have priority (allocate themselves before) over steps at higher depth:

$$\tilde{Q}^l_d(t+1) = (\tilde{Q}^l_d(t) + \sum_{(j,k)\text{-skills} \subset \mathcal{K}_l} A_{j,k}(t) - \sum_{m:S_m=\mathcal{K}_l} \lfloor \frac{h_m}{r_{1,1,1}} \rfloor U_m)^+. \tag{39}$$

Note that allocation of steps with skill requirements in $\mathcal{K}_l$ and $\mathcal{K}_{l'}$ for $l \neq l'$ are independent. Also, the agents with skills in $\mathcal{K}_l$ and $\mathcal{K}_{l'}$ for $l \neq l'$ are independent. So, if we define $\tilde{Q}_d(t)$:

$$\tilde{Q}_d(t+1) = (\tilde{Q}_d(t) + L \max_l (\sum_{(j,k)\text{-skills} \subset \mathcal{K}_l} A_{j,k}(t) - \sum_{m:S_m=\mathcal{K}_l} \lfloor \frac{h_m}{r_{1,1,1}} \rfloor U_m))^+, \tag{40}$$

then $\tilde{Q}_d(t)$ is a path-wise upper-bound on $\sum_l \tilde{Q}^l_d(t)$.

Consider the depth $d = 1$ first. Then we can follow the same steps as in the proof of Thm. 4 for the depth $d = 1$ case. Note that $\boldsymbol{\lambda} \in \alpha \mathcal{C}_{F,I}$ implies that for all l

$$\sum_{d=1}^{D} \sum_{j:(j,k)\text{-skill } \subset \mathcal{K}_l \cap \mathcal{A}_d} \lambda_j < (1-\alpha) \sum_{m:S_m=\mathcal{K}_l} \mu_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor,$$

and hence, for any $d \leq D$

$$\sum_{d'=1}^{d} \sum_{j:(j,k)\text{-skill } \subset \mathcal{K}_l \cap \mathcal{A}_{d'}} \lambda_j < (1-\alpha) \sum_{m:S_m=\mathcal{K}_l} \mu_m \lfloor \frac{h_m}{r_{1,1,1}} \rfloor.$$

This along with the same steps as in the proof of Thm. 4 gives that for some $\theta_1 > 0$ and $\theta_1 = \Omega(1)$, $\forall \theta < \theta_1$

$$\mathbf{E}[\exp(\theta_1 \tilde{Q}_1)] < \infty, \text{ for all } l.$$

Then, like the proof of Thm. 4 we can perform induction over d to prove that for some $\theta_D > 0$ and $\theta_D = \Omega(1)$, then for all $\theta < \theta_D$:

$$\mathbf{E}[\exp(\theta \tilde{Q}_D)] < \infty, \text{ for all } l,$$

where $\tilde{Q}_D(t)$ is an upper bound on $\sum_l Q_d^l(t)$, which is again the total number of unallocated steps in the system. We obtain the desired bound from this.

Induction from d to $d+1$ is similar to the proof of Thm. 4.